Failure in Safety-Critical Systems:

A HANDBOOK OF INCIDENT AND ACCIDENT REPORTING

Chris Johnson

Glasgow University Press

Glasgow University Press,

Publicity Services,
No. 2 The Square,
University of Glasgow,
Glasgow, G12 8QQ,
Scotland.

First printed October 2003.

ISBN 0-85261-784-4.

Contents

List of Figures

Preface

Incident reporting systems have been proposed as means of preserving safety in many industries, including aviation [308], chemical production [162], marine transportation [387], military acquisition [287] and operations [806], nuclear power production [382], railways [664] and healthcare [105]. Unfortunately, the lack of training material or other forms of guidance can make it very difficult for engineers and managers to set up and maintain reporting systems. These problems have been exacerbated by a proliferation of small-scale local initiatives, for example within individual departments in UK hospitals. This, in turn, has made it very difficult to collate national statistics for incidents within a single industry.

There are, of course, exceptions to this. For example, the Aviation Safety Reporting System (ASRS) has established national reporting procedures throughout the US aviation industry. Similarly, the UK Health and Safety Executive have supported national initiatives to gather data on Reportable Injuries, Diseases and Dangerous Occurrences (RIDDOR). In contrast to the local schemes, these national systems face problems of scale. It can become difficult to search databases of 500,000 records to determine whether similar incidents have occurred in the past.

This book, therefore, addresses two needs. The first is to provide engineers and managers with a practical guide on how to set up and maintain an incident reporting system. The second is to provide guidance on how to cope with the problems of scale that affect successful local and national incident reporting systems.

In 1999, I was asked to help draft guidelines for incident reporting in air traffic control throughout Europe. The problems of drafting these guidelines led directly to this book. I am, therefore, grateful to Gilles le Gallo and Martine Blaize of EUROCONTROL for helping me to focus on the problems of international incident reporting systems. Roger Bartlett, safety manager at the Maastricht upper air space Air Traffic Control center also provided valuable help during several stages in the writing of this book. Thanks are also due to Michael Holloway of NASA's Langley Research Center who encouraged me to analyze the mishap reporting procedures being developed within his organization. Mike O'Leary of British Airways and Neil Johnstone of Aer Lingus encouraged my early work on software development for incident reporting. Ludwig Benner, Peter Ladkin, Karsten Loer and Dmitri Zotov provided advice and critical guidance on the causal analysis sections. I would also like to thank Gordon Crick and Mark Bowell of the UK Health and Safety Executive, in particular, for their ideas on the future of national reporting systems.

I would like to thank the University of Glasgow for supporting the sabbatical that helped me to finish this work.

Chris Johnson, Glasgow, 2003.

Chapter 1

Abnormal Incidents

Every day we place our trust in a myriad of complex, heterogeneous systems. For the most part, we do this without ever explicitly considering that these systems might fail. This trust is largely based upon pragmatics. No individual is able to personally check that their food and drink is free from contamination, that their train is adequately maintained and protected by appropriate signalling equipment, that their domestic appliances continue to conform to the growing array of international safety regulations [278]. As a result we must place a degree of trust in the organisations who provide the services that we use and the products that we consume. We must also, indirectly, trust the regulatory framework that guides these organisations in their commercial practices. The behaviour of phobics provides us with a glimpse of what it might be like if we did not possess this trust. For instance, a fear of flying places us in a nineteenth century world in which it takes several days rather than a few hours to cross the Atlantic. The SS United States' record crossing took 3 days, 10 hours and 40 minutes in July 1952. Today, the scheduled crossings by Cunard's QEII now take approximately 6 days. In some senses, therefore, trust and profit are the primary lubricants of the modern world economy. Of course, this trust is implicit and may in some cases be viewed as a form of complicit ignorance. We do not usually pause to consider the regulatory processes that ensures our evening meal is free of contamination or that our destination airport is adequately equipped.

From time to time our trust is shaken by failures in the infrastructure that we depend upon [70]. These incidents and accidents force us to question the safety of the systems that surround us. We begin to consider whether the benefits provided by particular services and products justify the risks that they involve. For example, the Valujet accident claimed the lives of a DC-9's passengers and crew when it crashed after takeoff from Miami. National Transportation Safety Board (NTSB) investigators found that SabreTech employees had improperly labelled oxygen canisters that were carried on the flight. These cannisters created the necessary conditions for the fire, which in turn led to the crash. Prior to the accident, in the first quarter of 1996, Valujet reported a net income of $10.7 million. After the accident, in the final quarter of 1996, Valujet reported a loss of $20.6 million. These losses do not take into account the additional $262 million costs of settlements with the victims relatives.

The UK Nuclear Installations Inspectorate's report into the falsification of pellet diameter data in the MOX demonstration facility at Sellafield also illustrates the consequences of losing international confidence [641] In the wake of this document, Japan, Germany and Switzerland suspended their ships to and from the facility. The United States' government initiated a review of BNFL's participation in a £4.4bn contract to decommission former nuclear facilities. US Energy Secretary Bill Richardson sent a team to England to meet with British investigators. British Nuclear Fuel's issued a statement which stated that they had nothing to hide and were confident that the US Department of Energy would satisfy itself on this point [106].

The Channel Tunnel fire provides another example of the commercial consequences of such adverse events. In May 1997, the Channel Tunnel Safety Authority made 36 safety recommendations after finding that the fire had exposed weaknesses in underlying safety systems. Insufficient staff training had led to errors and delays in dealing with the fire. Eurotunnel, therefore, took steps to address these concerns by implementing the short-term recommendations and conducting further studies to consider those changes that involved longer-term infrastructure investment. However, the UK Consumer Association mirrored more general public anxiety when its representatives stated that it was 'still worried' about evacuation procedures and the non-segregation of passengers from cars on the tourist shuttle trains [97] Te fire closed the train link between the United Kingdom and France for approximately six months and served to exacerbate Eurotunnel's 1995 loss of £925 million.

This book introduces the many different incident reporting techniques that are intended to reduce the frequency and mitigate the consequences of accidents, such as those described in previous paragraphs. The intention is that by learning more from 'near misses' and minor incidents, these approaches can be used to avoid the losses associated with more serious mishaps. Similarly, if we can identify patterns of failure in these low consequence events we can also reduce the longer term costs associated with large numbers of minor mishaps. In order to justify why you should invest your time in reading the rest of this work it is important to provide some impression of the scale of the problems that we face. It is difficult to directly assess the negative impact that workplace accidents have upon safe and successful production [283]. Many low-criticality and 'near miss' events are not reported even though they incur significant cumulative costs. In spite of such caveats, it is possible to use epidemiological surveys and reports from national healthcare systems to assess the effects of incidents and accidents on worker welfare.

1.1 The Hazards

Employment brings with it numerous economic and health benefits. It can even improve our life expectancy over those of us who may be unfortunate enough not to find work. However, work exposes us to a range of occupational hazards. The World Health Organisation (WHO) estimate that there may be as many as 250 million occupational injuries each year, resulting in 330,000 fatalities [872]. If work-related diseases are included then this figure grows to 1.1 million deaths throughout the globe [873]. About the same number of people die from malaria each year. The following list summarises the main causes of occupational injury and disease.

- *Mechanical hazards.* Many workplace injuries occur because of poorly designed or poorly screened equipment. Others occur because people work on, or with, unsafe structures. Badly maintained tools also create hazards that may end in injury. Musculo-skeletal disorders and repetitive strain injury are now the main cause of work-related disability in most of the developed world. The consequent economic losses can be as much as 5% of the gross national product in some countries [872]. The Occupational Safety and Health Administration's (OSHA) ergonomics programme has argued that musculo-skeletal disorders are the most prevalent, expensive and preventable workplace injuries in the United States. They are estimated to cost $15 billion in workers' compensation costs each year. Other hazards of the working environment include noise, vibration, radiation, extremes of heat and cold.

- *Chemical Hazards.* Almost all industries involve exposure to chemical agents. The most obvious hazards arise from the intensive use of chemicals in the textile, healthcare, construction and manufacturing industries. However, people in most industries are exposed to cleaning chemicals. Others must handle petroleum derivatives and various fuel sources. Chemical hazards result in reproductive disorders, in various forms of cancer, respiratory problems and an increasing number of allergies. The WHO now ranks allergic skin diseases as one of the most prevalent occupational diseases [872]. These hazards can also lead to metal poisoning, damage to the central nervous system and liver problems caused by exposure to solvents and to various forms of pesticide poisoning.

- *Biological hazards.* A wide range of biological agents contribute to workplace diseases and infections. Viruses, bacteria, parasites, fungi, moulds and organic dusts affect many different industries. Healthcare workers are at some risk from tuberculosis infections, Hepatitis B and C as well as AIDS. For agricultural workers, the inhalation of grain dust can cause asthma and bronchitis. Grain dust also contains mould spores that, if inhaled, can cause fatal disease [321].

- *Psychological Hazards.* Absenteeism and reduced work performance are consequences of occupational stress. These problems have had an increasing impact over the last decade. In the United Kingdom, the cost to industry is estimated to be in excess of £6 billion with over 40 million working days lost each year [90]. There is considerable disagreement over the causes of such stress. People who work in the public sector or who are employed in the service industries seem to be most susceptible to psychological pressures from clients and customers. High workload, monotonous tasks, exposure to violence, isolated work have all been cited as contributory factors. The consequences include unstable personal relationships, sleep disturbances and depression. There can be physiological consequences including higher rates of coronary heart disease and hypertension. Post traumatic stress disorder is also increasingly recognised in workers who have been involved in, or witnessed, incidents and accidents.

This list describes some of the hazards that threaten workers' health and safety. Unfortunately, these items tell us little about the causes of these adverse events or about potential barriers. For example, OSHA report describes the way in which a sheet metal worker was injured by a mechanical hazard:

> "...employee #1 was working at station #18 (robot) of a Hitachi automatic welding line. She had been trained and was working on this line for about 2 months... The lifting arm then rises and a robot arm moves out from the operator's side of the welding line and performs its task. Then there is a few seconds delay between functions as the robot arm finishes welding, rises, returns to home and the lifting arm lowers to home, ready for the finished length of frame steel to move on and another to take it's place. During the course of this operation the welding line is shut down intermittently so that the welding tips on the robot arms can be lubricated, preventing material build up. This employee, without telling anyone else or shutting down the line, tried to perform the lubrication with the line still in automatic mode. She thought this could be done between the small amount of time it took all parts to complete their functions and return to home. The employee did not complete the task in time, as she had anticipated. Her right leg was located between the protruding rods on the lifting arm and the openings the rods rest in. Her leg was trapped. When other employees were alerted, they had trouble trying to switch the line to manual because the computer was trying to complete it's function and the lifting arm was trying to return to home. The result was that one employee used a crowbar to help relieve pressure on her leg and another used the cellenoid which enabled the lifting arm to rise. The employee received two puncture wounds in the thigh (requiring stitches) and abrasions to the lower leg. Management once again instructed employees working this line on the serious need to wait until all functions are complete, the line shut down and not in the automatic mode before attempting any maintenance." (OSHA Accident Report ID: 0352420).

It is possible to identify a number of factors that were intended to prevent this incident from occurring. Line management had trained the employees not to intervene until the robot welding cycle was complete. Lubrication was intended to be completed when the line was 'shut down' rather than in automated mode. It is also possible to identify potential factors that might have been changed to prevent the accident from occurring. For example, physical barriers might have been introduced into the working environment so that employees were prevented from intervening during automated operations. Similarly, established working practices may in some way have encouraged such risk taking as the report comments the management 'once again' instructed employees to wait until the line was shut down. These latent problems created the context in which the incident could occur [698]. The triggering event, or catalyst, was the employee's decision that she had enough time to lubricate the device. The lack of physical barriers then left her exposed to the potential hazard once she had decided to pursue this unsafe course of action. Observations about previously unsafe working practices in this operation may also have done little to dissuade her from this intervention.

Figure 1.1 provides a high level view of the ways in which incidents and accidents are caused by catalytic failures and weakened defences. The diagram on the left shows how the integration of working practices, working environment, line management and regulatory intervention together support a catalytic or triggering failure. Chapter 3 will provide a detailed analysis of the sources for such catalytic failures. For now, however, it is sufficient to observe that there a numerous potential causes ranging from human error through to stochastic equipment failures through to deliberate violations of regulations and working practices. It should also be apparent that there may be catalytic failures of such magnitude that it would be impossible for any combination of the existing structures to support, for any length of time. In contrast, the diagram on the right of Figure 1.1 is intended to illustrate how weaknesses in the integration of system components can increase an application's vulnerability to such catalytic failures. For example, management might strive to satisfy the requirements specified by a regulator but if those requirements are flawed then there is a danger that the system will be vulnerable to future incidents. These failures in the supporting infrastructure are liable to develop over a much longer timescale than the triggering events that place the system under more immediate stress.

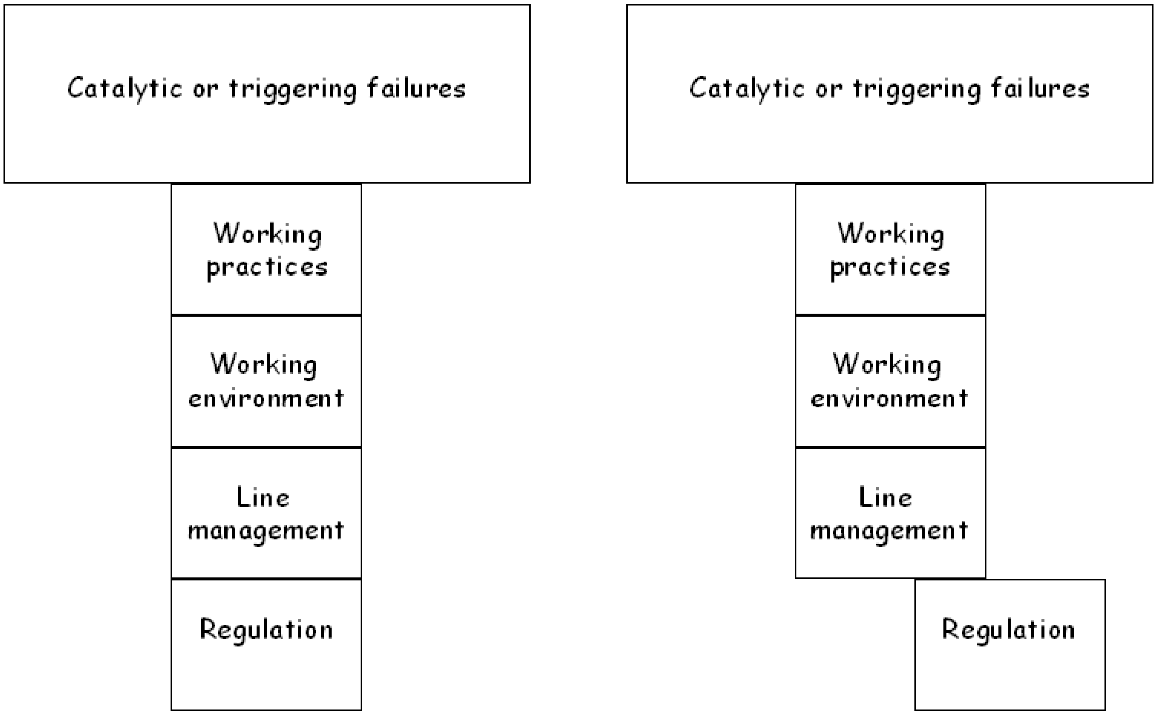

Figure 1.1: Components of Systems Failure

The diagrams in Figure 1.1 sketch out one view of the way in which specific failures place stress on the underlying defences that protect us from the hazards what were listed in previous paragraphs. A limitation of these sketches is that they provide an extremely static impression of a system as it is stressed by catalytic failures. In contrast, Figure 1.2 provides a more process oriented view of the development of an occurrence or critical incident. Initially, the systems is in a 'normal' state. Of course, this 'normal' state need not itself be safe if there are flaws in the working practices and procedures that govern everyday operation. The systems may survive through an incubation period in which any residual flaws are not exposed by catalytic failures. This phase represents a 'disaster waiting to happen'. However, at some point such an event does cause the onset of an incident or accident. These failures may, in turn, expose further flaws that trigger incidents elsewhere in the same system or in other interrelated applications. After the onset of a failure, protection equipment and other operators may intervene to mitigate any consequences. In some

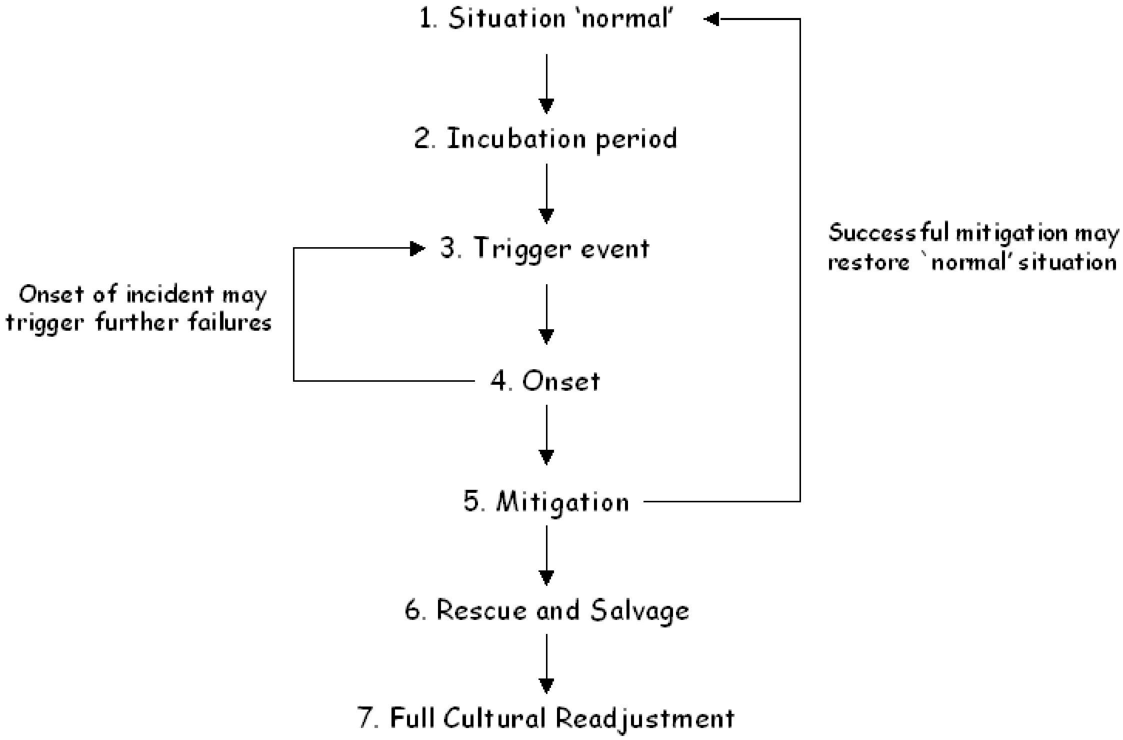
1. Situation 'normal'
2. Incubation period
3. Trigger event
4. Onset
5. Mitigation
6. Rescue and Salvage
7. Full Cultural Readjustment
Onset of incident may trigger further failures
Successful mitigation may restore 'normal' situation

Figure 1.2: Process of Systems Failure

cases, this may return the system to a nominal state in which no repair actions are taken. This has potentially dangerous implications because the flaws that were initially exposed by the triggering event may still reside in the system. Alternatively, a rescue and salvage period may be initiated in which previous shortcomings are addressed. In particular, a process of cultural readjustment is likely if the potential consequences of the failure have threatened the continued success of the organisation as a whole. For example, the following passage comes from a report that was submitted to the European Commission's Major Accident Reporting System (MARS) [229]:

> "At 15:30 the crankcase of an URACA horizontal action 3 throw pump, used to boost liquid ammonia pressure from 300 psi to 3,400 psi, was punctured by fragments of the failed pump-ram crankshaft. The two operators investigating the previously reported noises from the pump were engulfed in ammonia and immediately overcome by fumes. Once the pump crackcase was broken, nothing could be done to prevent the release of the contents of the surge drum (10 tonnes were released in the first three minutes). The supply of ammonia from the ring main could only be stopped by switching off the supply pump locally. No one were able to do this as the two gas-tight suits available were preferentially used for search and rescue operations, and thus release of ammonia continued. Ammonia fumes quickly began to enter the plant control room and the operators hardly had the time to sound the alarms and start the plant shut-down before they had to leave the building using 10 minutes escape breathing apparatus sets. During the search and rescue operation the fire authorities did not use the gas-tight suits and fumes entered the gaps around the face piece and caused injuries to 5 men. The ammonia cloud generated by the initial release drifted off-site and remained at a relatively low level." (MARS report 814).

A period of normal operation led to an incubation period in which the pump-ram crankshaft was beginning to fail and required maintenance. The trigger event involved the puncture of the pump's crankcase when the ram crankshaft eventually failed. This led to the onset of the incident in which two operators were immediately overcome. This then triggered a number of further, knock-on failures. For instance, the injuries to the firemen were caused because they did not use gas tight suits during their response to the initial incident. In this case, only minimal mitigation was possible as operators did not have the gas tight suits that were necessary in order to isolate the ammonia supply from the ring main. Those suits that were available were instead deployed to search and rescue operations.

Many of the stages shown in Figure 1.2 are based on Turner's model for the development of a system failure [790]. The previous figure introduces a mitigation phase that was not part of this earlier model. This is specifically distinguished from Turner's rescue and salvage stage because it reflects the way in which operators often intervene to 'cover up' a potential failure by taking immediate action to restore a nominal state. In many instances, individuals may not even be aware that such necessary intervention should be reported as a focus for potential safety improvements. As Leveson points out, human intervention routinely prevents the adverse consequences of many more occurrences than are ever recorded in accident and incident reports [486]. This also explains our introduction of a feedback loop between the mitigation and the situation normal phases. These features were not necessary in Turner's work because his focus was on accidents rather than incidents. Figure 1.2 also introduces a feedback loop between the onset and trigger phases. This is intended to capture the ways in which an initial failure can often have knock-on effects throughout a system. It is very important to capture these incidents because are increasingly common as we move to more tightly integrated, heterogeneous application processes.

Previous paragraphs have sketched a number of ways in which particular hazards contribute to occupational injuries. They have also introduce a number of high-level models that can be used to explain some of the complex ways in which background failures and triggering events combine to expose individuals to those hazards. The following sections build on this analysis by examining the likelihood of injury to individuals in particular countries and industries. We also look at the costs of these adverse events to individuals and also to particular industries. The intention is to reiterate the importance of detecting potential injuries and illnesses before they occur.

1.1.1 The Likelihood of Injury and Disease

Work-place incidents and accidents are relatively rare. In the United Kingdom, approximately 1 in every 200 workers reports an occupational illness or injury resulting in more than three days of absence from employment every year [331]. OSHA estimates that the rate of work-related injuries and illnesses dropped from 7.1 per year for every 100 workers in 1997 to 6.7 in 1998 [652]. These figures reflect significant improvements over the last decade. For example, the OSHA statistics show that the number of work-related fatalities has almost been halved since it was established by Congress in 1971. The Australian National Occupational Health and Safety Commission report that the rate of fatality, permanent disability or a temporary disability resulting in an absence from work of one week or more was 2.2 per 100 in 1997-8, 2.5 in 1996-7, 2.7 in 1995-6, 2.9 in 1994-95, 3.0 in 1993-4, 2.8 in 1992-3 [44]. The following figures provide the same data per million hours worked: 13 in 1997-8, 14 in 1996-7, 16 in 1995-6, 16 in 1994-5, 17 in 1993-4, 19 in 1992-3.

These statistics hide a variety of factors that continue to concern governments, regulators, managers, operators and the general public. The first cause for concern stems from demographic and structural changes in the workforce. Many countries continue to experience a rising number of workers. This is both due to an increasing population and to structural changes in the workforce, for instance increasing opportunities for women. In the United Kingdom, the 1% fall between 1998 and 1999 in the over 3 day injury rate is being offset by a (small) rise in the total number of injuries from 132,295 to 132,307 in 1999-2000 [331]. Similarly the OSHA figures for injury and illness rates show a 40 % decline since 1971. At the same time, however, U.S. employment has risen from 56 million workers at 3.5 million worksites to 105 million workers at nearly 6.9 million sites [652]. Population aging will also have an impact upon occupational injury statistics. Many industrialised countries are experiencing the twin effects of a falling birth rate and a rising life expectancy. This will increase pressure on the workforce for higher productivity and greater contributions to retirement provision. Recent estimates place the number of people aged 60 and over at 590 million worldwide.

By 2020, this number is projected to exceed 1,000 million [873]. Of this number, over 700 million older people will live in developing counties. These projections are not simply significant for the burdens that they will place on those in work. Older elements of the workforce are often the most likely to suffer fatal work-related injuries. In 1997-98, the highest rate of work-related fatalities in Australia occurred in the 55 plus age group with 1.3 deaths per 100 employees. They were followed by the 45-49 and 50-54 age groups with approximately 0.8 fatalities per 100 employees. The lowest number of fatalities occurred in workers under that age of 20 with 0.2 deaths per 100 employees. It can be difficult to interpret such statistics. For example, they seem to indicate that the rising risks associated with aging outweigh any beneficial effects from greater expertise across the workforce. Alternatively, the statistics may indicate that younger workers are more likely to survive injuries that would prove fatal to older colleagues. The UK rate of reportable injury is lower in men aged 16-19 than all age groups except for those above 55 [326]. However, the HSE report that the differences between age groups are not statistically significant when allowing for the higher accident rates for those occupations that are mainly performed by younger men. There is also data that contradicts the Australian experience. Young men, aged 16-24, face a 40% higher relative risk of all workplace injury than men aged 45-54 even after allowing for occupations and other job characteristics.

The calculation of health and safety statistics has also been effected by social and economic change. Part-time work has important effects on the calculation of health and safety statistics per head of the working population [652, 326]. The rate of injury typically increases with the amount of time exposed to a workplace risk. However, it is possible to normalise the rate using an average number of weekly hours of work. The rate of all workplace injury in the UK is 8.0 per 100 for people working less than 16 hours per week. For people working between 16 and 29 hours per week it is 4.3, between 30 and 49 hours it is 3.8, between 50 and 59 it is 3.2 and people working 60 or more hours per week have an accident rate of 3.0 per 100 workers per annum. People who work a relatively low number of hours have substantially higher rates of all workplace and reportable injury than those working longer hours. The relatively high risk in workers with low hours remains after allowing for different occupational characteristics [326]. The growth of temporary work has similar implications for some economies. In the UK, the rate of injury to workers in the first 6 months is double that of their colleagues who have worked for at least a year. This relatively high risk for new workers remains after allowing for occupations and hours of work. 57% temporary workers have been with their employer for less than 12 months.

Figure 1.1 shows that accident rates are not uniformly distributed across industry sectors. For example, the three day rate for agriculture and fishing in the United Kingdom is 1.2 per 100 employees. The same rate for the services industries is approximately 0.4 per 100 workers.

Industry	UK		Germany		France	Spain		Italy
	1993	1994	1993	1994	1993	1992	1993	1991
Agriculture	7.3	8.5	6.0	6.7	9.8	9.1	5.4	18.4
Utilities	0.5	0.6	3.1	4.3	5.6	12.5	10.1	4.4
Manufacturing	1.6	1.2	2.3	1.6	2.3	6.7	4.9	3.3
Construction	8.9	6.9	7.9	8.0	17.6	21.0	19.3	12.8
Transport	2.2	2.0	7.2	7.5	6.5	13.0	10.7	11.2
Other services	0.3	0.4	1.0	1.2	1.9	1.4	1.5	0.9
All Industries	1.2	0.9	3.3	3.2	3.9	6.4	5.1	5.5

Table 1.1: Industry Fatality Rates in UK, Germany, France, Spain & Italy [324]

Accidents rates also different with gender. Positive employment practices are exposing increasing numbers of women to a greater variety of risks in the workplace. The overall Australian National Occupational Health and Safety Commission rate of 2.2 injuries and illnesses per 100 workers hides a considerable variance [44]. For males the rate was 2.9 per 100 workers whilst it was 1.3 for females.

In 1997-8, the industries with the highest number of male fatalities were Transport and Storage (66) and Manufacturing (64), while for females Accommodation, Cafes and Restaurants (4) and Property and Business Services (4) were the highest. The male fatalities were mainly employed as Plant and Machine Operators, and Drivers (91). Female fatalities were mainly employed as Managers and Administrators (5). These differences may decline with underlying changes in workplace demographics. However, UK statistics suggest some significant residual differences between the genders:

> "the rate of all workplace injury is over 75% higher in men than women, reflecting that men tend to be employed in higher risk occupations. After allowing for job characteristics, the relative risk of workplace injury is 20% higher in men compared with women. Job characteristics explain much of the higher rate of injury in men but not all because men still have an unexplained 20% higher relative risk". [326]

Table 1.1 illustrates how the rate of industrial injuries differs within Europe. Such differences are more marked when comparisons are extended throughout the globe. However, it is not always possible to find comparable data:

> "The evaluation of the global burden of occupational diseases and injuries is difficult. Reliable information for most developing countries is scarce, mainly due to serious limitations in the diagnosis of occupational illnesses and in the reporting systems. WHO estimates that in Latin America, for example, only between 1 and 4% of all occupational diseases are reported. Even in industrialised countries, the reporting systems are sometimes fragmented." [873]

For example, the Australian statistics cited in previous paragraphs include some cases of coronary failure that would not have been included within the UK statistics. These problems are further exacerbated by the way in which local practices affect the completion of death certifications and other reporting instruments. For instance, the death of a worker might have been indirectly caused by a long running coronary disease or by the immediate physical exertion that brings on a heart attack. It is important to emphasise that even if it were possible to implement a consistent global reporting system for workplace injuries, it would still not be possible to directly draw inferences about the number of incidents and accidents directly from that data. Many incidents still go unreported even if well-established reporting systems are available. A further limitation is that injury and fatality statistics tell us little or nothing about 'near miss' incidents that narrowly avoided physical harm.

1.1.2 The Costs of Failure

In 1996 the UK Health and Safety Executive estimated that workers and their families lost approximately £558 million per year in reduced income and additional expenditure from work-related injury and ill health [322]. They also estimated that the loss of welfare in the form of pain, grief and suffering to employees and their families was equivalent to a further £5.5 billion. These personal costs also have wider implications for employers, for the local economy and ultimately for national prosperity. The same study estimated that the direct cost to employers was approximately £2.5 billion a year; £0.9 billion for injuries and £1.6 billion for illness. In addition, the loss caused by avoidable accidental events that do not lead to injury was estimated at between £1.4 billion and £4.5 billion per year. This represents 4-8% of all UK industrial and commercial companies' gross trading profits.

Employers also incur costs through regulatory intervention. These actions are intended to ensure that a disregard for health and safety will be punished whether or not an incident has occurred. Tables 1.2 and 1.3 summarise the penalties imposed by United States' Federal and State inspectors in the fiscal year 1999 [652]. Regulatory actions imposed a cost of $151,361,442 beyond the immediate financial losses incurred from incidents and accidents. These figures do not account for the numerous competitive disadvantages that are incurred when organisations are associated with high-profile failures [675].

Violations	Percent	Type	Penalties
646	0.8	Willful	$24,460,318
50,567	66	Serious	$50,668,509
1,816	2	Repeat	$8,291,014
226	0.3	Failure to abate	$1,205,063
408	0.01	Unclassified	$3,740,082
23,533	30	Other	$1,722,338
77,196	Total		$90,087,324

Table 1.2: Federal Inspections Fiscal Year 1999

Violations	Percent	Type	Penalties
441	0.3	Willful	$12,406,050
57,010	40	Serious	$35,441,267
2,162	1.5	Repeat	$4,326,620
785	0.5	Failure to abate	$2,860,972
46	0.0002	Unclassified	$2,607,900
82,120	40	Other	$3,631,309
202,962	Total		$61,274,118

Table 1.3: State Inspections Fiscal Year 1999

1.2 Social and Organisational Influences

These statistics illustrate the likelihood and consequences of occupational injuries. It is important, however, to emphasise that this data suffers from a number of biases. Many of the organisations that are responsible for collaring the statistics are also responsible for ensuring that mishap frequencies are reduced over time. Problems of under-reporting can also complicate the interpretation of national figures. There is often a fear that some form of blame will attach itself to those organisations that return an occupational health reporting form. The OSHA record keeping guidelines stress that:

> "Recording an injury or illness under the OSHA system does not necessarily imply that management was at fault, that the worker was at fault, that a violation of an OSHA standard has occurred, or that the injury or illness is compensable under workers' compensation or other systems." [653]

However, in many counties including the United States, organisations that have a higher reported rate of occupational illness or injury become the focus of increasing levels of regulatory inspection and intervention. This has a certain irony because, as OSHA acknowledge, relatively low levels of reported injuries and illnesses may be an indicator of poor health and safety management:

> "...during the initial phases of identifying and correcting hazards and implementing a safety and health program an employer may find that its reported rate increases. This may occur because, as an employer improves its program, worker awareness and thus reporting of injuries and illnesses may increase. Over time, however, the employer's ... rate should decline if the employer has put into place an effective program." [648]

It is instructive to examine how our analysis relates to previous work on enhancing the safety of hazardous technologies. Two schools of thought can be identified; the first stems from the 'normal accident' work of Perrow [675]; the second stems from the idea of 'high reliability' organisations [718].

1.2.1 Normal Accidents?

Perrow argues that the characteristics of high-risk technologies make accidents inevitable, in spite of the effectiveness of conventional safety devices. These characteristics include complexity and tight coupling. Complexity arises from our limited understanding of some transformation stages in modern processing industries. It stems from complex feedback loops in systems that rely on multiple, interacting controls. Complexity also stems from many common-mode interconnections between subsystems that cannot easily be isolated. More complex systems produce unexpected interactions and so can provoke incidents that are harder to rectify.

Perrow also argues that tight coupling plays a greater role in the adverse consequences of many accidents than the complexity of modern technological systems. This arises because many applications are deliberately designed with narrow safety margins. For example, a tightly coupled system may only permit one method of achieving a goal. Access to additional equipment, raw materials and personnel is often limited. Any buffers and redundancy that are allowed in the system are deliberately designed only to meet a few specified contingencies. In contrast, Perrow argues that accidents can be avoided through loose coupling. This provides the time, resources and alternative paths to cope with a disturbance.

There is evidence to contradict parts of Perrow's argument [710, 684]. Some 'high reliability' organisations do seem to be able to sustain relatively low incident rates in spirit of operating complex processes. Viller [847] identifies a number of key features that contribute to the perceived success of these organisations:

- The leadership in an organisation places a high priority on safety.
- High levels of redundancy exist even under external pressures to trim budgets.
- Authority and responsibility are decentralised and key individuals can intervene to tackle potential incidents. These actions are supported by continuous training and by organisational support for the maintenance of an appropriate safety culture.
- Organisational learning takes place through a variety of means, including trial and error but also through simulation and hypothesis testing.

These characteristics illustrate the important role that incident reporting plays for 'high reliability' organisations. Such applications are an important means of supporting organisational learning. Table 1.4 summarises the main features of 'Normal Accident' theory and 'High Reliability' organisations. Sagan [718] used both of these approaches to analyse the history of nuclear weapons safety. His conclusions lend weight to Perrow's pessimistic assessment that some accidents are inevitable. They are significant because they hold important implications for the interpretation both of incident and accident reports. For example, Sagan argues that much of the evidence put forward to support high reliability organisations is based on data that those organisations help to produce. Accounts of good safety records in military installations are often dependent on data supplied by the military. This is an important caveat to consider during the following pages in which we will present incident and accident statistics. We may not always be able to rely upon the accuracy of information that organisations use to publicise improvements in their own safety record. Sagan also argues that social pressures act as brakes on organisational learning. He identifies ways in which stories about previous failures have been altered and falsified. He then goes on to show how the persuasive effects of such pressures can help to convince the originators of such stories that they are, in fact, truthful accounts of incidents and accidents. This reaches extremes when failures are re-painted as notable successes.

1.2.2 The Culture of Incident Reporting

Sagan's work shows that a variety of factors can affect whether or not adverse events are investigated. Thes factors affect both individuals and groups within safety-critical organisations. The impact of cultural influences, of social and legal obligations, cannot be assessed without regard to individual differences. Chapter 3 will describe how subjective attitudes to risk taking and to the violation of rules can have a profound impact upon our behaviour. For now it is sufficient to observe that each of the following influences will affect individuals in a number of different ways.

In some groups, it can be disloyal to admit that either you or your colleagues have made a mistake or have been involved in a 'failure'. These concerns take a number of complex forms. For example,

High Reliability Organisations	Normal Accidents Theory
Accidents can be prevented through good organisational design and management	Accidents are inevitable in complex and tightly coupled systems.
Safety is the priority organisational objective.	Safety is one of a number of competing objectives.
Redundancy enhances safety: duplication and overlap cam make a reliable system out of unreliable parts.	Redundancy often causes accidents: it creates interactive complexity and encourages risk taking.
Decentralised decision-making is needed to permit prompt and flexible operating responses to surprises	De-centralised control is needed for complex systems but centralised control is needed for tight coupling.
A culture of reliability enhances safety by encouraging uniform and appropriate responses by operators	A military model of intense discipline and isolation is incompatible with democratic values
Continuous operations, training and simulations can create and maintain high reliability operations.	Organisations cannot train for unimagined, highly dangerous or politically unpalatable operations
Trial and error learning from accidents can be effective and can be supplemented by anticipation and simulations	Denial of responsibility, faulty reporting and reconstruction of history cripples learning efforts.

Table 1.4: Competing Perspectives on Safety with Hazardous Technologies [718]

individuals may be prepared to report failures. However, individuals may be reluctant to face the retribution of their colleagues should their identity become known. These fears are compounded if they do not trust the reporting organisation to ensure their anonymity. For this reason, NASA go to great lengths to publicise the rules that protect the identity of contributors to the US Aviation Safety Reporting System.

Companies can support a good 'safety culture' by investing in and publicising workplace reporting systems. A number of factors can, however, undermine these initiatives. The more active a company is in seeking out information about previous failures then the worse its safety record may appear. It can also be difficult to sustain the employee protection that encourages contributions when incidents have economic as well as safety implications. Individuals can be offered re-training after a first violation, re-employment may be required after a second or third.

The social influence of a company's ‘safety culture’ is reinforced by the legal framework that governs particular industries. This is most apparent in the regulations that govern what should and what should not be reported to national safety agencies. For example, the OSHA regulations follow Part 1904.12(c) of the Code of Federal Regulations. These require that employers record information about every occupational death; every nonfatal occupational illness; and those nonfatal occupational injuries which involve one or more of the following: loss of consciousness, restriction of work or motion, transfer to another job, or medical treatment (other than first aid). [653] As we shall see, this focus on accidents rather than ‘near-miss’ incidents reflects an ongoing debate about the scope of Federal regulation and enforcement in the United States.

It is often argued that individuals will not contribute to reporting systems unless they are protected from self-incrimination through a ‘no blame’ policy [700]. It is difficult for organisations to preserve this ‘no blame’ approach if the information that they receive can subsequently be used during prosecutions. Conversely, a local culture of non-reporting can be reinforced or instigated by a fear of legal retribution if incidents are disclosed. These general concerns characterise a range of more detailed institutional arrangements. For example, some European Air Traffic Management providers operate under a legal system in which all incidents must be reported to the police. In neighbouring countries, the same incidents are investigated by the service providers themselves and, typically, fall under an informal non-prosecution agreement with state attorneys. Other countries have more complex legal situations in which specific industry arrangements also fall under more general regional and national legislation. For example, the Utah Public Officers and Employees' Ethics Act and the Illinois' Whistle Blower Protection Act are among a number of state instruments that have been passed to protect respondents. These local Acts provide for cases that are also covered by Federal statutes including the Federal False Claims Act or industry specific provision for Whistle Blowers such as section 405 of the Surface Transportation Assistance Act. This has created some disagreement about whether state legislation preempts federal law in this area; several cases have been conducted in which claimants have filed both common law and statutory suits at the same time. Cases in Texas and Minnesota have shown that Federal statutes provide a base-line and not a ceiling for protection in certain states. Such legal complexity can deter potential contributors to reporting systems.

There are other ways in which the legislative environment can affect reporting behaviour. For example, freedom of information and disclosure laws are increasing public access to the data that organisations can hold. The relatives or representatives of people involved in an accident can potentially use these laws to gain access to information about previous incidents. In such circumstances, there is an opportunity for punitive damages to be sought if previous, similar incidents were reported but not acted upon. These concerns arose in the aftermath of the 1998 Tobacco Settlement with cigarette manufacturers in the United States. Prior to this settlement, states alleged that companies had conspired to withhold information about the adverse health effects of tobacco [580].

The legislative environment for accident and incident reporting is partly shaped by higher-level political and social concerns. For example, both developed and developing nations have sought to deregulate many of their industries in an attempt to encourage growth and competition. Recent initiatives to liberalise the Indian economy have highlighted this conflict between the need to secure economic development whilst also coordinating health and safety policy. The Central Labour Institute has developed national standards for the reporting of major accidents. However, the Directorate General of Factory Advice Services and the Labour Institutes have not developed similar guidelines for incident and occurrence reporting. The focus has been on developing education and training programmes that can target specific health and safety issues after industries have become established within a region [156].

Some occupational health and safety reporting system have, however, been extended to explicitly collect data about both actual accidents and 'near-miss' incidents. For example, employers in the UK are guided by the Reporting of Injuries, Diseases and Dangerous Occurrences Regulations (RIDDOR) 1995. These cover accidents which result in an employee or a self-employed person dying, suffering a major injury, or being absent from work or unable to do their normal duties for more than three days. They also cover 'dangerous occurrences' that do not result in injury but have the potential to do significant harm [320]. These include:

- The collapse, overturning or failure of load-bearing parts of lifts and lifting equipment.
- The accidental release of a biological agent likely to cause severe human illness.
- The accidental release of any substance which may damage health.
- The explosion, collapse or bursting of any closed vessel or associated pipework.
- An electrical short circuit or overload causing fire or explosion.
- An explosion or fire causing suspension of normal work for over 24 hours.

Similarly, Singapore's Ministry of Manpower requires that both accidents and 'dangerous occurrences' must be reported. Under the fourth schedule of the national Factory Act, these may 'under other circumstances' have resulted in injury or death [742]. The detailed support that accompanies the act provide exhaustive guidance on the definition of such dangerous occurrences. These are taken to include incidents that involve bursting of a revolving vessel, wheel, grindstone or grinding wheel.

Dangerous occurrences also range from electrical short circuit or failure of electrical machinery, plant or apparatus, attended by explosion or fire or causing structural damage to an explosion or failure of structure of a steam boiler, or of a cast-iron vulcaniser.

A duty to report on incidents and accidents does not always imply that information about these occurrences will be successfully acted upon. This concern is at the heart of continuing attempts to impose a 'duty to investigate' upon UK employers. At present, the UK regulatory framework is one in which formal accident investigation of the most serious incidents is undertaken by specially trained investigators. Employers are not, in general, obliged to actively finding out what caused something to go wrong. Concern about this situation led to a 1998 discussion document that was published by the Health and Safety Commission (HSC). It was observed that:

> "At present, there is no law which explicitly requires employers to investigate the causes of workplace accidents. Many employers do undertake accident investigation when there has been an event in the workplace which has caused injury in order to ensure lessons are learnt, and although there is no explicit legal duty to investigate accidents there are duties under some health and safety law which may lead employers to undertake investigation. The objective of a duty to investigate accidents would be to ensure employers draw any appropriate lessons from them in the interests of taking action to prevent recurrence." [314]

There are many organisational reasons why a body such as the HSC would support such an initiative. The first is the face-value argument that such a duty to investigate accidents and incidents would encourage employers to adopt a more pro-active approach to safety. The second is that such a duty would help to focus finite regulatory resources by following the deregulation initiated in the UK under the Robens Committee [709]. This group responded to the mass of complex regulations that had emerged from the plethora of nineteenth century factory acts. As industries merged and emerged, it was difficult for employers to know which parts of each act actually applied to their business. As a result, the Robens Committee helped to propose what became the Health and Safety at Work Act (1974). Key sections of the Roben report [701] argued that:

> "We need a more effective self-regulating system... It calls for better systems of safety organisation, for more management initiatives, and for more involvement of work people themselves. The objectives of future policy must, therefore, include not only increasing the effectiveness of the state's contribution to safety and health at work but also, and more importantly, creating conditions for more effective self-regulation" [709]

The same concerns over the need to target finite regulatory resources and the need to encourage pro-active intervention by other organisations also inspired attempts in the United States to establish OSHA's Cooperative Compliance Programme. This focused on the 12,000 employers that had the highest reported mishap rates. Those companies that agreed to participate and invest in safety management programs were to be offered a reduced likelihood of OSHA inspection. This was estimated to be a reduction from an absolute certainty of inspection down to approximately 30% [648]. This policy was intended to leverage OSHA resources by encouraging commercial investment in safety. It was also intended to provide OSHA with a means of targeting finite inspection resources. However, employers' organisations claimed that it introduced new roles and responsibilities for the Federal organisation. The US Chamber of Commerce helped to present a case before the US Court of Appeals that succeeded in blocking OSHA's plans. The Assistant Secretary of Labour for Occupational Safety and Health argued:

> "The goal of Cooperative Compliance Programme (CCP) is to use OSHA's limited resources to identify dangerous work sites and work in partnership with management and labour to find and fix hazards. America's taxpayers expect nothing less for their continued support and funding of OSHA. This lawsuit is frivolous; it has no merit and aims only to hinder our ability to protect working men and women from often life-threatening hazards. The CCP is an enforcement program–not a regulation. We are confident that our program is lawful. Attempts by the National Association of Manufacturers and the

> U.S. Chamber of Commerce to throw-up legal roadblocks will only ensure that the most dangerous work sites in America remain that way, putting untold numbers of workers at risk." [397]

The CCP provides important insights into the regulatory environment in the United States. As a result of the legal action, OSHA was forced to build less formal partnerships with employers' organisations. The CCP is also instructive because OSHA produced detailed guidance on those measures that high-reporting organisations ought to introduce in order to address previous failures. Table 1.5 presents OSHA's guidelines [648] on how to assess the quality of accident investigation within an organisation. As can be seen, the investigation of 'near-miss' incidents, or occurrences in HSE terms, characterises an organisation at the highest level of safety management.

1	No investigation of accidents, injuries, near misses, or other incidents is conducted.
2	Some investigation of incidents takes place, but root cause may not be identified, and correction may be inconsistent. Supervisors prepare injury reports for lost time cases.
3	OSHA-101 (report form) is completed for all recordable incidents. Reports are generally prepared with cause identification and corrective measures prescribed.
4	OSHA-recordable incidents are always investigated, and effective prevention is implemented. Reports and recommendations are available to workers. Quality and completeness of investigations are systematically reviewed by trained safety personnel.
5	All loss-producing accidents and near-misses are investigated for root causes by teams or individuals that include trained safety personnel and workers.

Table 1.5: OSHA Levels of Accident and Incident Investigation

Different reporting systems have different definitions of what should and what should not be reported. These distinctions reflect national and international agreements about the nature of incidents and accidents. For instance, Table 1.6 embodies International Civil Aviation Organisation (ICAO) and EUROCONTROL requirements for incident and accident reporting in Air Traffic Control. As can be seen, this covers both specific safety-related incidents such as the loss of control in flight and also failures to provide adequate air traffic management services.

Table 1.6 provides domain dependent definitions of incidents and accidents. Each row provides explicit examples of occurrences in Air Traffic Management. It could not easily be used in the chemical or healthcare industries. It can still be difficult to apply these consequence based definitions of ATM incidents and accidents. For example, a loss of separation might be avoided if air crews spot each other and respond appropriately. Such an occurrence might be given a relatively low criticality assessment; no loss of separation occurred. However, it can also be argued that this incident ought to be treated *as if* an air proximity violation had occurred because air traffic control did not intervene to prevent it from happening. This approach is exploited within some European ATM service providers.

Further problems complicate the use of consequence based definitions of accidents and incidents, such as those illustrated in Table 1.6. Individuals may not be able to observe the consequences of the adverse events that they witness. For example, maintenance teams are often remote from the operational outcomes of their actions. As a result, organisations such as the UK Civil Aviation Authority approve specific lists of occurrences that must be reported. For instance, the Ground Occurrence Report Form E1022 is used for the notification of defects found during work on aircraft or aircraft components which are considered worthy of special attention [10]. In contrast to Table 1.6, the following list includes procedural errors and violations, such as incorrect assembly, as well as

Occurrence	Category	Definitions of an Occurrence
Accidents	Mandatory	Mid-air collision, controlled flight into terrain, ground collision between aircraft, ground collision between aircraft and obstruction. Other accidents of special interest including loss of control in flight due to VORTEX or meteorological conditions.
Incidents	Mandatory	Loss of air separation, near controlled flight into terrain, runway incursion, inability to provide ATM services, breach in ATM system security.
Other occurrences	Voluntary	Anything which has serious safety implications but which is neither an accident nor an incident.

Table 1.6: Distinctions between Accidents and Incidents in Air Traffic Control

observations of potential component failure, such as overheating of primary or secondary structure:

- Defects in aircraft structure such as cracks in primary or secondary structure, structural corrosion or deformation greater than expected
- Failures or damage likely to weaken attachments of major structural items including flying controls, landing gear, power plants, windows, doors, galleys, seats and heavy items of equipment
- When any component part of the aircraft is missing, believed to have become detached in flight
- Overheating of primary or secondary structure
- Incorrect assembly
- Failure of any emergency equipment that would prevent or seriously impair its use
- Critical failures or malfunction of equipment used to test aircraft systems or aircraft units
- Any other occurrence or defect considered to require such notification.

The ICAO list of air traffic incidents relied upon an analysis of the potential consequences of any failure. In contrast, the CAA definition of ground maintenance incidents was built from a list of errors, violations and observations of potential failures. These differences can be explained in terms of the intended purpose of these definitions. In the former case, the list of ATM accidents and incidents was intended as a guideline for safety managers in national service providers. They are assumed to have the necessary investigative resources, analytical insights and reconstruction capabilities to assess potential outcomes once incidents have been reported. However, the CAA reporting procedures provide direct guidance for maintenance personnel. These individuals are not expected to anticipate the many different potential outcomes that can stem from the failures that they observe. Such criticality assessments must be performed by the line management who receive and interpret the information from incident reporting systems. These differences illustrate the difficulty of developing a priori definitions of accidents and incidents that ignore the purpose to which those definitions will be put.

Some authors have constructed more general definitions of accidents and incidents. For instance, Perrow [675] proceeds by distinguishing between four levels of any system. Unlike most regulatory definitions, such as that illustrated in Table 1.6, Perrow does not focus directly on the likely consequences of a failure but rather looks at those portions of a system that were effected by an incident or accident:

1. *parts.* The first level of any system represent the smallest components that are likely to be considered during an accident investigation. They might include objects such as a valve.

2. *unit.* These are functionally related collections of parts. For example, a motor unit is built from several individual component parts.

3. *subsystem.* These are composed from individual units. For example, the secondary coolant system of a nuclear reactor will contain a steam generator and a water return unit.

4. *the plant or system.* This is the highest level involved in an accident. Beyond this it is only profitable to think in terms of the impact of an accident on the environment.

In Perrow's terms, accidents only involve those failures that affect levels three and four of this hierarchy. Incidents disrupt components at levels one and two. This definition is critical for the normal accidents argument that Perrow proposes in his book. He argues that 'engineered safety functions' cannot reliably be constructed to prevent some incidents from becoming accidents at levels three and four. Unfortunately, however, these distinctions raise a number of problems for our purposes. Definitions of incidents and accidents must serve the pragmatic role of helping individual workers to know what should, and what should not, be reported. It is unclear whether people would ever be able to make the distinctions between levels 2 and 3 that would be required under this scheme.

There are further practical problems in applying such structural distinctions between accidents and incidents. As with consequential definitions, it may be difficult for any individual to determine the scope of any failure as it occurs. They may fail to realise that the failure of a level one valve will create knock-on effects that compromise an entire level four system. The social and cultural issues that were introduced in previous sections also affect the interpretation of accidents and incidents. This is illustrated in Figure 1.3. From the viewpoint of person A, the system is operating 'abnormally' as soon as it moves from state 1 to state 2. Person B holds different beliefs about what is, and what is not, normal. As a result, they only consider that an incident has occurred when the system moves from state 2 to state 3. The different viewpoints shown in this sketch can arise for a number of reasons. For example, Person A may have been trained to identify the transition between states 1

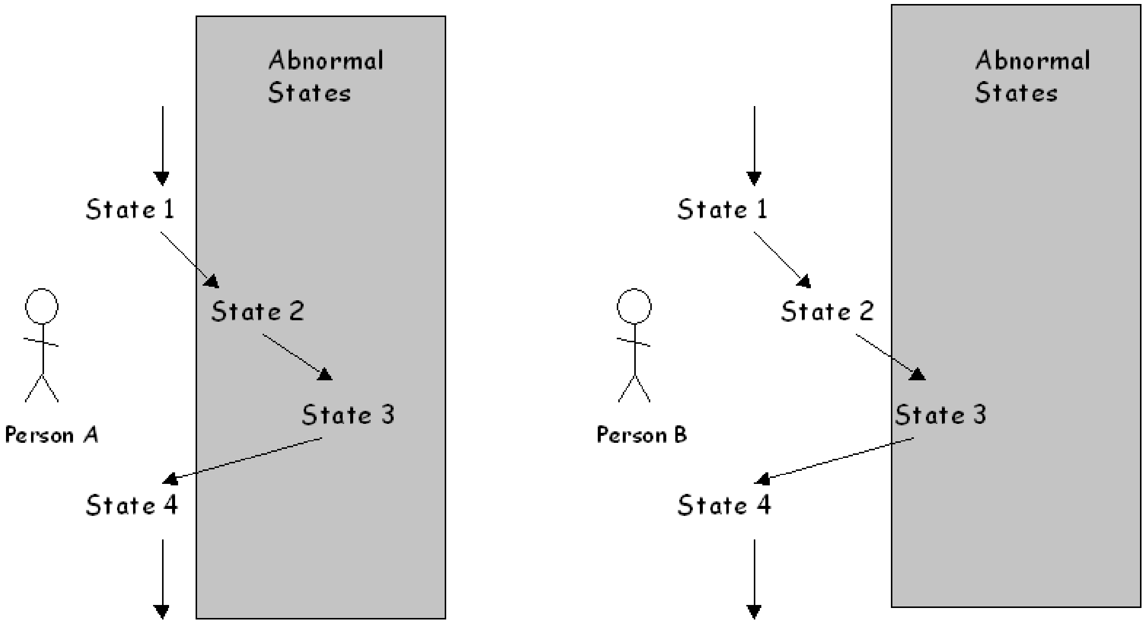

Figure 1.3: Normal and Abnormal States

and 2 as potentially hazardous. Alternatively, person B may exhibit individual attitudes to risk that can dispose them not to report hazardous incidents. Figure 1.3 can also illustrate how attitudes to hazards may change over time. For example, the figure on the left might represent an initial attitude when the system is initially installed. Over time, dangerous working practices can become the norm. It can be difficult for many individuals to question such established working practices even if they violate recognised rules and regulations. Over time, these dangerous practices may themselves become sanctioned by procedures and regulations. This is illustrated by what Diane Vaughan has called "normalised deviance" in the events leading to the Challenger accident [846]. Under such circumstances, the figure on the right might represent the prevailing view of normal and abnormal states.

The previous analysis helps to identify a number of possible approaches to the definition of what an incident actually is. These can be summarised as follows:

- *open definitions.* This approach encourages personnel to report any failure as a safety-related incident. It is exploited by the Air Navigation Services Division of the Swedish Civil Aviation Administration. As a result they receive several thousand reports per year ranging from the failure of lights or heating systems through to potential air proximity violations. The open approach to the definition of an incident avoids some of the problems with more restricted definitions, see below. However, it can also lead to a dilution of the safety reporting system with more general concerns. In the Air Navigation Services Division this approach is well supported by trained 'Gatekeepers' who filter low priority reports from more serious occurrences. The entire system is, however, dependent on the skill and insight of these personnel and their ability to perform a timely analysis of the initial reports.

- *closed definitions.* Closed systems lie at the other extreme from open definitions such as those exploited by the Swedish Air Traffic Control organisation. These systems provide rigid definitions or enumerations for those incidents that are to be reported to the system. All staff are trained to recognise these high priority occurrences and all other incidents are handled through alternative mechanisms. The difficulty with this approach is that the introduction of new equipment can have a profound impact upon the sorts of incidents that will occur. As a result, these enumerations must be revised over time. Otherwise, staff will not report new incidents but will instead continue to wait for occurrences that are now prevented by more secure defences.

- *consequential definitions.* These represent a subset of the closed approach, described above. Incidents and accidents are distinguished either by their actual outcomes or by the *probable worst case* consequences. For example, the US Army regulations distinguish between class A to D accidents whose consequences range from $1,000,000 or more (class A) to between $2,000 and $10,000 (class D) [806] Class E incidents result in less that $2,000 damage but interrupt an operational or maintenance mission. Class F incidents relate to Foreign Object Damage and are restricted to aviation operations. As we have seen, the problem here is that it can be difficult for operators to predict the possible consequences of a failure without further investigation and analysis. As a result, these definitions tend to be applied by investigators and analysts after an initial warning or report has been generated.

- *structural definitions.* This is a further example of a closed approach which has strong links to consequential definitions. The consequences of a failure are assessed for each of several layers of a system. Incidents affect the lower level components whilst accidents involve the system as a whole. There are a number of practical problems in applying this as a guide for incident reporting. there are also theoretical problems when individual component faults may cause a fatality, for example through electrocution, even though the system as a whole continues to satisfy its functional requirements. A strict interpretation of such events would rank them as an incident and not an accident in Perrow's terms [675].

- *procedural definitions.* This is another example of a closed approach. Rather than focusing on the anticipated outcome of a failure, procedural definitions look at violation of the prescribed methods. The problem here is that the individuals who witness violations may fail to recognise them as violations, especially if they have become part of standard working practices. Such problems also affect incident reporting systems that ask operators to comment on 'anything unusual'.

- *pragmatic definitions.* Some incident reporting systems take a particularly pragmatic approach to the definition of what should and what should not be reported. They are often characterised by the phrase 'target the doable'. This characterises systems that have been established within larger organisations that may not, as a whole, support the recommendations of the scheme. Some of the pioneering attempts to establish incident reporting systems within the UK National Health Service deliberately focused on those occurrences that individual consultants felt that they could address; incidents stemming from wider acquisitions policy or even from other clinical departments were deliberately excluded.

- *special issues.* Finally, some incident reporting systems deliberately focus on key issues. For instance, the European Turbulent Wake incident reporting system was established with help from the UK Meteorological Service in response to concerns about a number of occurrences involving commercial flights. Other systems are deliberately focused to elicit information from key personnel who may be under-represented in existing incident databases. For example, schemes have been initiated to encourage incident reporting from General Aviation and military pilots rather than commercial pilots. Other schemes have focused on eliciting information from medical and surgical staff rather than nursing personnel.

The preceding discussion should illustrates the difficulty of providing a single definition of accidents and incidents. These problems stem from the different ways in which different people must use these definitions. The person witnessing an adverse occurrence must know whether or not it is worthwhile reporting. Safety managers may apply different criteria when determining whether or not an incident report merits a full-scale investigation or whether it can be dealt with at a more local level. National authorities may apply further criteria when deciding whether national trends indicate a need for regulatory intervention.

It is important to emphasise that the distinction between an incident and an accident is not firm and cannot be made a priori. The same set of events may be reclassified at several stages in the investigation and analysis of an occurrence. These must not be arbitrary decisions. Later chapters will stress the need to provide a documented justification for such changing assessments. However, there are often important pragmatic reasons for such actions. For example, a number of European air traffic control agencies have not reported any major accidents in recent years. As a result, some air traffic service providers have begun to treat certain 'critical incidents' as-if they were accidents, even though no loss of life or property has occurred. The intention is to rehearse internal procedures for dealing with more critical events when, and if, they do occur. Such decisions also focus attention and resources on the causes of these incidents. To summarise, simple distinctions between accidents and incidents ignore the underlying complexity that characterises the ways in which different national and international organisations treat technological failure. Different definitions are used, and may indeed be necessary, to support different stages of an organisation's response to incidents and accidents.

1.3 Summary

It is difficult to estimate the costs when human error, systems failure or managerial weakness threatens safety. Employers face a number of direct costs when their employees are injured. The UK Health and Safety Executive estimate that occupational injuries cost employers around 4-8% of their gross trading profit; currently approximately £6 billion. There are also indirect costs that accrue when regulators intervene. In the United States Federal and State inspectors levied penalties for health and safety violations that totalled $151,361,442 for the fiscal year 1999. Incident or occurrence reporting systems enable companies to identify potential failures before they occur. They provide insights that can be used to guide risk assessment during subsequent development.

Incident reporting systems provide regulators with data that can be used to guide any necessary intervention. They help to prioritise health and safety initiatives and awareness raising campaigns. They can also be used to address public concerns, for example the creation of a national incident reporting system for UK railways followed shortly after the Ladbroke Grove and Southall accidents. At an international level, incident reporting systems provide means of ensuring that lessons are effectively shared across national boundaries. The following chapter introduced the challenges that must be addressed if these claimed benefits are to be realised.

Chapter 2

Motivations for Incident Reporting

This chapter explains why many organisations develop incident reporting systems. The intention is often to identify potential failures before an accident occurs. The higher frequency of less critical mishaps and near-miss events also supports statistical analysis that cannot reliably be performed on relatively infrequent accidents. Data and lessons from one system can be shared with the operators of other similar applications. The following pages also identify limitations that are often forgotten by the proponents of incident reporting systems. Many submissions do little more than remind their operators of hazards that are well understood but are difficult to avoid. The resources used by a reporting system might alternatively fund safety improvements. Managers of successful reporting systems can be overwhelmed by a mass of data about relatively trivial mishaps. Later sections go on to review issues of confidentiality and scope that help to determine whether the claimed benefits outweigh the perceived costs of operating these systems.

2.1 The Strengths of Incident Reporting

The US Academy of Science recommended that a nationwide mandatory reporting system should be established to improve patient safety [453]. They argued that this system should initially be based around hospitals but that eventually other ‘care settings’ should be included. The International Civil Aviation Organisation has published detailed guidance on the manner in which reporting systems must be implemented within signatory states [384].

> “(The assembly) urges contracting states to undertake every effort to enhance accident prevention measures, particularly in the areas of personnel training, information feedback and analysis and to implement voluntary and non-punitive reporting systems, so as to meet the new challenges in managing flight safety, posed by the anticipated growth and complexity of civil aviation”.
> (Resolution A31-10: Improving accident prevention in civil aviation)
>
> “(The assembly) urges all Contracting States to ensure that their aircraft operators, providers of air navigation services and equipment, and maintenance organisations have the necessary procedures and policies for voluntary reporting of events that could affect aviation safety” (ICAO Resolution A32-15: ICAO Global Aviation Safety Plan)

The US Coast Guard and the Maritime Administration have helped to establish a voluntary international maritime information safety system. This is intended to receive, analyse, and disseminate information about unsafe occurrences. They argue that these ‘non-accidents’ or ‘problem events’ provide an untapped source of data. They can be used as indicators of safety-levels in the maritime community and provide the information necessary to prevent accidents before they happen [830]. The goals of the system are to reduce the frequency of marine casualties, to reduce the extent of injuries and property damage (including environmental damage), and to create a safer and more efficient shipping transportation system and mariner work environment.

The Council of the European Union had similar concerns when it drafted the 1996 directive on the control of major accident hazards. This has become more widely known as the Sveso II directive; it was named after the town in Italy where 2,000 people had to be treated following a release of tetrachlorodibenzoparadioxin (Dioxin) in 1976:

> “Whereas, in order to provide for an information exchange and to prevent future accidents of a similar nature, Member States should forward information to the Commission regarding major accidents occurring in their territory, so that the Commission can analyse the hazards involved, and operate a system for the distribution of information concerning, in particular, major accidents and the lessons learned from them; whereas this information exchange should also cover ‘near misses’ which Member States regard as being of particular technical interest for preventing major accidents and limiting their consequences.” [187]

The Transportation Safety Board of Canada [622] identified a number of reasons to justify the creation of its own confidential incident reporting system. They argued that incident data will support the Board’s studies on a wide range of safety-related matters including operating procedures, training, human performance and equipment suitability. The analysis of incident reports can also help to identify widespread safety deficiencies that might not have been detected from individual reports submitted to regional centres. Greater insights into national and international transportation safety issues can be gained by collating accident/incident reports and by comparing it with data from other agencies.

These individual initiatives across a range of industries illustrate the increasing importance of incident reporting within safety management systems [444]. They can also be used to identify common arguments that justify the development and maintenance of incident reporting systems:

1. Incident reports help to find out why accidents DONT occur. Many incident reporting forms identify the barriers that prevent adverse situations from developing into a major accident. These insights help analysts to strengthen those safeguards that have already proven to be effective barriers in ‘near miss’ incidents.

2. The higher frequency of incidents permits quantitative analysis. It can be argued that many accidents stem from atypical situations. They, therefore, provide relatively little information about the nature of future failures. In contrast, the higher frequency of incidents provides greater insights into the relative proportions of particular classes of human ‘error’, systems ‘failure’, regulatory ‘weakness’ etc.

3. They provide a reminder of hazards. Incident reports provide a means of monitoring potential problems as they recur during the lifetime of an application. The documentation of these problems increases the likelihood that recurrent failures will be noticed and acted upon.

4. Feedback keeps staff ‘in the loop’. Incident reporting schemes provide a means of encouraging staff participation in safety improvement. In a well-run system, they can see that their concerns are treated seriously and are acted upon by the organisation. Many reporting systems also produce newsletters that can be used to increase awareness about regional and national safety issues.

5. Data (and lessons) can be shared. Incident reporting systems provide the raw data for comparisons both within and between industries. If common causes of incidents can be observed then, it is argued, common solutions can be found. However, in practice, the lack of national and international standards for incident reporting prevents designers and managers from gaining a clear view of the relative priorities of such safety improvements.

6. Incident reporting schemes are cheaper than the costs of an accident. The relatively low costs of managing an incident reporting scheme should be offset against the costs of failing to prevent an accident. This is a persuasive argument. However, there is also a concern that punitive damages may be levied if an organisation fails to act upon the causes of an incident that subsequently contribute towards an accident.

7. May be required to do it. The final argument in favour of incident reporting is that these schemes are increasingly being required by regulatory agencies as evidence of an appropriate safety culture. This point is illustrated by the ICAO resolutions A31-10 and A32-15 and by the EC Seveso II directive that were cited on previous pages.

Many of these arguments require little additional explanation. For example, it it sufficient to cite the relevant ICAO resolutions to demonstrate that member states should implement incident reporting systems. However, some of these apparent justifications for incident reporting are more controversial. For example, we have argued that the higher number of incidents can be used to drive statistical analyses of the problems that lead to a far smaller number of accidents. Heinrich's [340] pioneering studies in occupational health and safety suggested an approximate ratio of one accident to thirty occurrences involving major injuries to three hundred 'near-miss' incidents. More recently, Bird [84] proposed a ratio of one accident, involving serious or disabling injuries, to ten minor injuries to 30 incidents involving property damage to six hundred incidents resulting in no visible damage. He based this on a statistical analysis of 1.5 million reported incidents. The work of Heinrich, Bird and their colleagues have led to the 'Iceberg' model of incident data. Any accident is the pinnacle, or more properly the nadir, of a far larger number of incidents. The consequences of this form of analysis seem clear. Incident reports provide a far richer data sources for organisational learning and the 'control' of major accidents.

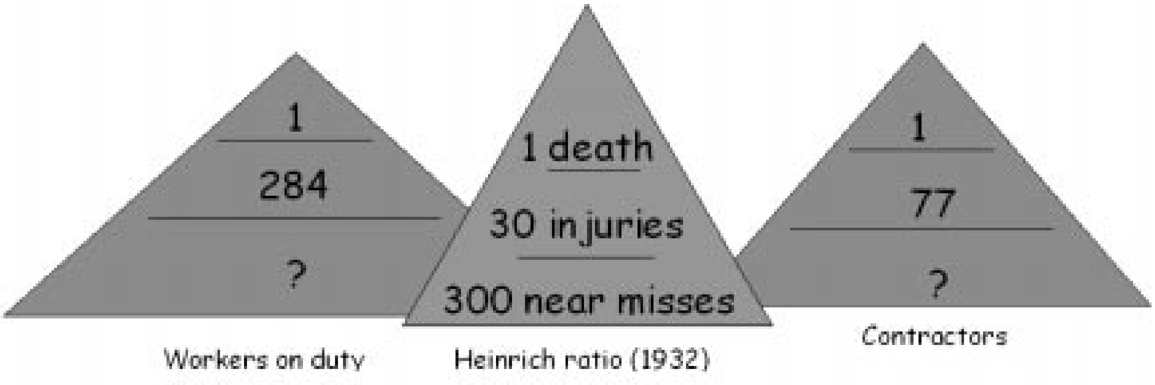

Figure 2.1: Federal Railroad Administration Safety Iceberg

Figure 2.1 illustrates a number of caveats that can be made about the Iceberg model. The central pyramid represents the results of Heinrich's initial study. On either side, the diagram presents the proportion of fatal to non-fatal injuries reported for different groups of workers in the US rail system based on Federal Railway Administration data from 1997 to 2000. Direct railroad employees or 'workers on duty' suffered a total of 119 fatalities and 33,738 injuries. Contractors experienced 31 fatalities and 1,466 injuries in the same period. The first problem is that the FRA has no reliable means of calculating the number of 'near miss' incidents over this period. As a result, it is only possible to examine the relationship between fatal work related deaths and injuries. Workers had a Heinrich ratio of one fatality for every two hundred and eighty-four injuries. The ratio for contractors was one fatality to seventy-seven injuries.

Further problems arise when we interpret these ratios. They might show that contractors are less likely to be injured than 'workers on duty'. An alternate way of expressing this is to say that contract staff are more likely to be killed than injured when compared to other employees. However, these ratios provide a very impoverished measure of probability. They do not capture the comparative risk exposure of either group. For example, the smaller number of fatal accidents to contractors may stem from a proportionately smaller number of workers. Contract workers are more likely than full-time, direct staff to be involved in high-severity incidents [874]. Alternatively, it can be argued that contractors are more reluctant to report work-related injuries than 'directly' employed staff.

This line of analysis is important because it questions the reliability of the data that can be obtained to calculate Heinrich ratios.

The argument that statistical data about incidents can be used to predict potential accidents is based on the premise that incidents are accidents in the making. It is assumed that incidents share the same root causes as more serious occurrences Van der Schaaf [843, 840] provides preliminary data from the Dutch chemical industry to confirm this premise. Glauz, Bauer and Migletz [291] also found a correlation between traffic conflicts and accidents. Other have exploited a more qualitative approach by looking for common contributory factors in both incidents and accidents. For instance, Helmreich, Butler, Taggart, and Wilhelm [341] have attempted to show that poor Crew Resource Management (CRM) causes both incidents and accidents. They then use this analysis to propose a predictive Cockpit Management Attitudes Questionnaire that can assess individual attitudes towards crew communication, coordination, and leadership issues.

A great deal of safety-related research rests on the assumption that incidents are good predictors of potential accidents. Wright has recently challenged this view in her statistical analysis of Scottish railways Confidential Incident Reporting and Analysis System (CIRAS) [874]. This confidential system elicits information about less 'critical' incidents. All accidents must, in contrast, be reported to a specialist unit within the UK Health and Safety Executive. Her work, therefore, focuses on 'near misses' and unsafe acts near the base of the Iceberg model. A near-miss has the potential to lead to a more serious occurrence, for example:

> "A Driver overshot a station platform by one and a half coach lengths. The Driver experiences wheelslip which may have been due to rail contamination. This did not result in any damage or injury" [874]

An unsafe act occurs when operator intervention actively undermines the safety of their system:

> "A Driver stated that when requested by the Signaller to do a controlled stop to assess railhead conditions he carries out this procedure assuming exceptional conditions i.e., reduced speed rather than normal speed. A controlled stop test carried out in this manner would not indicate the braking capacity in normal conditions and lead to an incorrect assumption that normal working may be resumed" [874]

Wright was able to conduct follow-up interviews with the staff who had submitted a confidential form from a total collection of 165 reports. A causal analysis was conducted using guidelines in the systems classification handbook and was validated by inter-rater reliability trials [197]. Occurrences were first assessed to identify technical and human factors issues. If a human factors 'failure' was identified then it was categorised as either proximal, distal or intermediate. Proximal factors include a range of human failures at the 'sharp end'. Intermediate factors relate to training or communications failures between high-level management and front-line staff. Distal factors relate to organisational and managerial issues that are remote from the workplace. Table 2.1 provides a comparison of the high-level causes of the 'near misses' and unsafe acts. The discrepancy between the number of reports and the total number of causal factors in this table can be explained by the fact that an incident can involve one or more causal factors.

Category	Near Miss (total 155)	Unsafe Acts (total 223)
Technical	20.7% (32)	1.3% (3)
Proximal	27.7% (43)	23.3% (52)
Intermediate	21.9% (34)	21.2% (47)
Distal	29.7% (46)	54.3% (121)

Table 2.1: Causal Comparison of CIRAS Incidents and Unsafe Acts

As can be seen, technical faults and failures seem to occur more frequently in near miss events than in unsafe acts. Conversely, distal factors such as organisation and managerial problems seem to occur more frequently as causal factors in unsafe acts. From this it follows that any analysis of 'near miss' events might fail to predict probable causes of actual incidents at the lower levels of the Iceberg model. These results can be explained in terms of the particular application area that Wright was studying. For example, near misses typically involved a failure to halt a train within the specified distance from a particular signal. These were often attributed to technical problems such as contaminated railheads. Unsafe acts were, in contrast, associated with the violations of company rules and procedures that govern driver behaviour on the UK railways. More work is required to confirm Wright's more general hypothesis that adverse events at the lower levels of the Iceberg model may provide poor predictors of accidents at the higher levels.

2.2 The Weaknesses of Incident Reporting

The most obvious limitation of incident reporting systems is that they can be expensive both to set up and to maintain. For instance, Leape notes that the Aviation Safety Reporting System spends about $3 million annually to analyse approximately 30,000 reports. This equates to about $100 (£66) per case.

> "These 'near miss' situations are far simpler to analyse than actual accidents, thorough investigation of which would almost certainly cost far more. It would be interesting to know, for example, the cost per case of investigations reported to the confidential enquiries. However, if we applied the figure from the Aviation Safety Reporting System to the 850,000 adverse events that are estimated to occur annually in the UK National Health Service, the cost of investigation would be £50 million annually." [480]

For comparison, it has been estimated that the cost of clinical negligence to health authorities and NHS Trusts was approximately £200 million in 1995-1996. The NHS summarised accounts for 1996-2001 include provision totalling £80 million with contingent liabilities of £1.6 billion [89]. Even when incident reporting systems are successfully established and maintained, a number of problems can limit their effectiveness. For instance, there is in reality very little sharing of incident data. For example, the European Confidential Aviation Safety Reporting Network ran between 1992 and 1999 with funding from the European Community. The network was intended to improve safety by passing on incident information to the aviation community. However, it was forced to close through lack of support from some sectors of the European aviation industry.

Further problems limit the transfer of incident information between organisations within an industry. For instance, Boeing operate an extensive system for collecting information about maintenance problems in their aircraft. They have successfully encouraged the exchange of data with airline operators. Unfortunately, however, there has been little coordination between airlines and groups of airlines about the format that this data should take. These formats are proprietary in the sense that they have been tailored to meet the specific needs of the operating companies. As a result when Boeing attempt to collate the data that is being shared they must face the considerable task of translating between each of these different formats. Any conclusions that are drawn from this data must also account for the different reporting cultures and reporting practices that exist within different operating groups [472].

Incident reporting systems may also fail to keep staff 'in the loop'. Occasionally these systems develop into grandiose initiatives that fulfill the organisational ambitions of their proponents rather than directly addressing key safety issues. There is also a danger that incident reporting systems degenerate into reminders of failures that everyone knows exists but few people have the political or organisational incentives to address [409]. Similarly, they may recommend short-term fixes or expedients that fail to address the underlying causes of incidents. This is illustrated by the following report from NASA's Aviation Safety Reporting System (ASRS):

> "Problem: on landing, gear was unlocked but up. Contributing factors: busy cockpit. [I] did not notice the gear down-and-locked light was not on. Discovered: Gear up was discovered on landing. Corrective action: [I] was unable to hear gear warning horn because of new noise cancelling headsets. I recommend removal of one ear-piece in

> landing phase of flight to audible warning devices to be heard by pilot. The noise-cancelling headsets were tested by three people on the ground and all three noted that with the headsets active that the gear warning horn was completely masked by the headsets." [62]

This illustrates the strengths and weaknesses of many incident report schemes. They provide first-hand insights into operational problems. They can also provide pragmatic remedies to the challenges that poorly designed equipment creates. However, there is also a danger that immediate remedies to individual incidents will fail to address the root cause of a problem. The noise-correcting headphones were clearly not fit for purpose. The proposed remedy of removing one headphone provides a short-term fix for individual pilots. However, it does little to address the underlying problems for future product development.

Further problems limit the ways in which data can be shared between incident reporting schemes. Although some organisations have successfully exchanged information about the frequency of particular occurrences, there have been few attempts to ensure any consistency in their response to those incidents. This creates particular problems for the maritime and aviation industries where operators may read of different recommendations being made in different countries. The following excerpt comes from the Confidential Human Factors Incident Reporting Programme (CHIRP). CHIRP is the UK equivalent of the ASRS that was cited in the previous quotation. This excerpt offers a slightly different perspective on the problems of ambient noise in the cockpit:

> "Fortunately, I have no incident to report. I would like, however, to highlight a common practice by some airlines, including my employer, which I feel is a significant risk to flight safety: namely the practice of not using flight deck intercom systems in favour of half wearing a headset over one ear for VHF comms, whilst using the other ear, unaided, for cockpit communications. And all this in what are often not so quiet flight decks.
>
> I cannot believe that we do not hear much better with two ears than with one, and many are the times when I, and other colleagues of mine, have had to ask for the other crew member to repeat things because of aircraft noise in one ear, and ATC in the other with the volume turned high enough not to miss a call. Not the best answer in a busy terminal area after a long flight, and an unnecessary increase in stress factors. Myself and others have raised this point several times to our training and safety departments, all of which has fallen, pardon the pun, onto deaf ears. The stock answer is that there is no written down SOP on intercoms, and common agreed practice rules. In reality, the guy in the right hand seat has no influence without things getting silly.
>
> As even single ear-piece headsets are not incompatible with intercoms, I would have thought a compromise would be mandatory use of full headset and intercom at the busy times, say below a given flight level, with the option for personal preferences in the cruise. Volumes for different communication channels could be adjusted to suit, and surrounding noise significantly reduced. This would preclude the need to speak louder than usual to be heard, to ask for repetitions, and general ly improve the working environment. After all, if the CAA and other agencies have made intercoms mandatory in transport aircraft, it will be for a reason.
>
> CHIRP Comment: The use of headsets for the purpose of effective reception of RTF/intercom messages between flight crew members is not mandated. The certification requirement for an intercom system is to provide communication between all crew members in an emergency. The partial/full use of a headset in normal operations should be dependent on the ambient noise level on the flight deck. For this reason, some operators specify the headset policy by aircraft type and phase of flight, as the reporter suggests. [175]"

The US ASRS article, cited above, argues that only one headset should be used during landing in order to help the crew hear cockpit warnings. In contrast, the CHIRP report condemns this practice as a threat to flight safety. This apparent contradiction is resolved by the second report, which argues that the partial or full use of headsets should be determined by the level of ambient noise. However, this distinction is not made explicit in the first report. Such differences illustrate the inconsistencies that can arise between national incident reporting systems. They are also indicative of a need to improve communication between these systems if we are to achieve the benefits that are claimed for the exchange of incident data. The ASRS and CHIRP systems are run by 'not for profit' organisations. The problems of data exchange are many times worse when companies may yield competitive advantage through the disclosure of incident information.

Incident reporting systems can provide important reminders about potential hazards. However, in extreme cases these reminders can seem more like glib repetitions of training procedures rather than pro-active safety recommendations. This problem is compounded by the tendency to simply remind staff of their failures rather than to address the root causes, such as poor design or 'error inducing environments' [362]. Over time the continued repetition of these reminder statements from incident reporting systems is symptomatic of deeper problems in the systems that users must operate:

> "On pre-flight check I loaded the Flight Management Computer (FMC), with longitude WEST instead of EAST. Somehow the FMC accepted it (it should have refused it three times). During taxi I noticed that something was wrong, as I could not see the initial route and runway on the navigation map display, but I got distracted by ATC. After we were airborne, the senior cabin attendant came to the flight deck to tell us the cabin monitor (which shows the route on a screen to passengers) showed us in the Canaries instead of the Western Mediterranean! We continued the flight on raw data only to find out that the Heading was wrong by about 30-40 degrees. With a ceiling of 1,000 ft at our destination I could not wait to be on 'terra firma'. Now I always check the Latitude/Longitude three times on initialisation!"
>
> (Editorial note) A simple but effective safeguard against 'finger trouble' of the type described is for the pilot who does not enter the data to confirm that the information that he/she sees displayed is that which he/she would expect. Then, and only then, should the 'Execute' function button be pressed." [176]

The CHIRP feedback is well intended. It also reiterates recommended practices that have formed part of Crew/Cockpit Resource Management (CRM) training for almost twenty years [410]. UK Aeronautical Information Circular (AIC) 143/1993 (Pink) states that all crew must have completed an approved CRM course before January 1995. Joint Airworthiness Requirement Operational Requirements (JAR OPS) sub-part N, 1.945(a)(10) and 1.955(b)(6) and 1.965(e) extended similar requirements to all signatory states during 1998. There is a considerable body of human factors research that points to the dangers of any reliance on such reminders [699]. Effectiveness declines with each repetition that is made. It is depressing, therefore, that such data-entry problems continue to be a frequent topic in aviation reporting systems. These incidents are seldom the result of deliberate violations or aircrew negligence. They illustrate the usability problems that persist within Commercial Aviation and which cannot simply be 'fixed' by training in cockpit coordination [410].

Incident reporting systems must go beyond repeated reminders to be 'careful' if they are to preserve the confidence of those who contribute to them. The US ASRS recognise this by issuing two different forms of feedback in response to the reports that they receive. The Callback bulletin describes short-term fixes to immediate problems. In contrast, the DirectLine journal addresses more systemic causes of adverse events and 'near miss' incidents even if it has a more limited audience than its sister publication. For instance, the following excerpt is taken from a DirectLine analysis of the causes of several mishaps involving Pre-Departure Clearances:

> "The type of confusion experienced by this flight crew over their (Pre-Departure Clearance) PDC routing is potentially hazardous, as noted by a controller reporter to ASRS: 'It has been my experience ... that several times per shift aircraft which have received PDCs with amended routings, have not picked up the amendment ... I have myself on numerous occasions had to have those aircraft make some very big turns to achieve sep-

> aration.' (ACN # 233622). The sources consulted by ASRS suggested several potential solutions to this problem:
>
> - Standardise PDC formats, so that pilots will know where to look for routing information and revisions.
> - Show only one clearance line in a PDC, and insert any revisions into the clearance line. Make the revision section more visible by tagging it ('REVISION') or highlighting with asterisks or other eye -catching notation (*****).
> - Provide flight crews with training in how to recognise PDC revisions." [56]

There are limits to the safety improvements that can be triggered through initiatives in publications such as DirectLine. Some mishaps can only be addressed through industry cooperation and regulatory intervention. Others require international agreements. For example, reporting systems have had a limited impact on workload in aviation. Similarly, usability problems continue to affect new generations of computer systems for airline operations. Data entry in flight management systems continues to be error prone many years after the problem was first identified. These 'wicked problems' must be considered when ambitious proposals are made to extend aviation reporting into healthcare and other transportation modes.

2.3 Different Forms of Reporting Systems

There are several different types of reporting system. This section explains why concerns over retribution have led to anonymous and confidential schemes. It also explains how both national and local systems have been set up to ensure that recommendations do not simply degenerate into reminders about known problems.

2.3.1 Open, Confidential or Anonymous?

The FAA launched the Global Aviation Information Network (GAIN) initiative as an attempt to encourage national and commercial organisations to exchange occurrence data. The Office of System Safety that drove the GAIN proposal within the FAA identified four main barriers to the success of such a system. These can be summarised as follows:

> "1. Punishment/Enforcement. First, potential information providers may be concerned that company management and/or regulatory authorities might use the information for punitive or enforcement purposes. In the US, significant progress has been made on this issue. Following the example of the UK, the FAA issued a policy statement in 1998 to the effect that information collected by airlines in their Flight Operations Quality Assurance (FOQA) programs, in which flight data recorder information is collected routinely, will not ordinarily be used against the airlines or pilots for enforcement purposes. In January 2000, the US President announced the creation of the Aviation Safety Action Programme (ASAP), in which airlines will collect reports from pilots, mechanics, dispatchers, and others about potential safety concerns, and made a commitment analogous to the FOQA commitment not to use the information for enforcement purposes. In April 2000, Congress enacted legislation that requires the FAA to issue a rule to develop procedures to protect air carriers and their employees from enforcement actions for violations that are discovered from voluntary reporting programs, such as FOQA and ASAP programs.
>
> 2. Public Access. Another problem in some countries is public access, including media access, to information that is held by government agencies in certain countries. This problem does not affect the ability of the aviation community to create GAIN, but it could affect the ability of government agencies in some countries to receive information from GAIN. Thus, in 1996 the FAA obtained legislation that requires the agency to protect voluntarily supplied aviation safety information from public disclosure. This

> will not deprive the public of any information to which it would otherwise have access, because the agency would not otherwise receive the information; but on the other hand, there is a significant public benefit for the FAA to have the information because it helps the FAA prevent accidents and incidents. The FAA is now developing regulations to implement that legislation...
>
> 3. Criminal Sanctions. A problem in some countries is the fear of criminal prosecution for regulatory infractions. Such a fear would be an obvious obstacle to the flow of aviation safety information. This has not historically been a major problem in the U.S., but the trend from some recent accidents is troubling.
>
> 4. Civil Litigation. Probably the most significant problem, certainly in the U.S., is the concern that the information will be used against the reporter in accident litigation. Some have suggested that, as was done in relation to the public disclosure issue, the FAA should seek legislation from Congress to protect aviation safety information from disclosure in litigation. In comparison with the public disclosure issue, however, the chances of obtaining such legislation are probably very remote; and a failed attempt to obtain such legislation could exacerbate the situation further because these disclosure issues are now determined in court, case by case, and a judge who is considering this issue might conclude that a court should not give protection that Congress refused to give." [308]

Incident reporting systems have addressed these concerns in a number of different ways. For instance, it is possible to identify three different disclosure policies. Anonymous systems enable contributors to entirely hide their identity. Confidential systems allow the limited disclosure of identity but only to trusted parties. Finally, open systems reveal the identity of all contributors. The impact of the distinctions between open, confidential and anonymous systems cannot be under-emphasised. In anonymous systems, contributors may have greater confidence in their submission; safe in the knowledge that they can avoid potential 'retribution'. However there is a danger that spurious reports will be filed. This problem is exacerbated by the fact that it is difficult to substantiate anonymous reports to determine whether they really did occur in the manner described. Investigators cannot simply ask about an incident within a workgroup without the possibility of implicating the contributor. This would remove the protection of confidentiality and could destroy the trust that is fundamental to the success of such systems. The distinctions between open, anonymous and confidential systems are also blurred in many existing applications. For example, the Swedish Air Traffic Control organisation (Luftfartsverket Flygtrafikjänsten) encourages the open contribution of incident reports. However, normal reporting procedures direct submissions through line supervisors. There is a danger that this might dissuade contributions about the performance of these supervisors. As a result, procedures exist for the confidential submission of incident reports via more senior personnel.

Trust and Technological Innovation

Distinctions between confidential, anonymous and open systems are intended to sustain the confidence and *trust* of potential participants. In a confidential system, contributors trust that only 'responsible' parties will receive identification information. The implications of this for the operation of any reporting system are illustrated by the approach taken with the CIRAS system that covers UK railways. This receives paper-based forms from train drivers, maintenance engineers and other rail staff. A limited number of investigators are responsible for processing these forms. They will conduct follow-up interviews in-person or over the telephone. These calls are not made to the contributor's workplace for obvious reasons. The original report form is then returned to the employee. No copies are made. Investigators type up a record of the incident and conduct a preliminary analysis. However, all identifying information is removed from the report before it is submitted for further analysis. From this point it is impossible to link a particular report to a particular employee. The records are held on a non-networked and 'protected' data base. This data itself is not revealed to industry management. However, anonymized reports are provided to managers every three months.

Incident reporting systems increasingly rely on computer-based applications . The Swedish Air Traffic Control system, mentioned above, is an example of this. Controllers in airfields in the more remote areas of Northern Sweden can receive rapid feedback on a report using this technology. However, electronic submission creates a number of novel and complex challenges for systems that attempt to preserve anonymity. These concerns are illustrated by the assurances that are provided to contributors on the Swiss Anaesthesia Critical Incident Reporting System. These include a commitment that they 'will NOT save any technical data on the individual reports: no E-mail address and no IP-number (a number that accompanies each submitted document on the net)' [755]. The use of computer-based technology not only raises security problems in the maintenance of trust during the transmission and storage of electronic documents, it also offers new and more flexible ways of maintaining incident reporting systems. For example, the US Department of Energy's Computerised Accident/Incident Reporting System (CAIRS) exploits an access control mechanism to tailor the level of confidentiality that is afforded to particular readers of particular incident reports. The CAIRS database is used to collect and analyse reports of injuries, illnesses, and other accidents that are submitted to the Department of Energy by their staff or contractors. The following paragraphs provide a brief overview of the innovative way in which the confidentiality of information is tied to particular access rights.

> "When you are granted access to CAIRS, you will be assigned an organisational jurisdiction. This jurisdiction may be for a specific organisation or for a complete contractor, area office, or field office. This jurisdiction assignment will determine the records that will be selected when the default organisation selection is utilised in many of the reports and logs. The default can be over-ridden by entering the desired organisation codes in the appropriate input boxes.
>
> CAIRS reports contain personal identifiers (names and social security numbers) and information regarding personal injury or illness. In order to prevent an unwarranted invasion of personal privacy, all personal identifiers are masked from the view of general users whenever any logs or reports are generated.
>
> The default registration for CAIRS does not provide access to any privacy information. If you require access to privacy information in order to perform your job function, you may apply for access to that information." [655]

It can be difficult to communicate the implications of such computer-based security measures to non-computer literate employees. There is a natural reluctance to believe in the integrity of such safeguards given continuing press coverage about the vulnerability of 'secure' systems [1]. The ability to access this data over the web might compound such misgivings.

Workplace Retribution and Legal Sanction

At least two different classes of problems exist in more open systems. Later paragraphs will address the issues that arise when trying to integrate a pro-active safety culture into a punitive legal system. There is a natural reluctance to implicate oneself or one's colleagues when subsequent investigations might directly threaten their livelihood and wellbeing. The second set of problems arise from a justified fear of persecution from colleagues or employers. These fears are natural if, for example, the subject of a report is a person in a position of authority or if the report reflects badly upon such a person. These individuals are likely to have a strong influence upon the career prospects and promotion opportunities of their more junior colleagues. The long term consequences of any actual or implied criticism can be extremely serious. Such concerns have long been apparent in the 'cockpit gradient'; co-pilots have extreme difficulty in challenging even minor mistakes made by a Captain. Co-Pilots have been known to remain silent even when their colleague's behaviour threatened the lives of everyone on board [733].

There are other reasons why individuals can be reluctant to contribute to incident reporting systems. There may be a fatalism that such an individual or group will suppress the report. If the report focuses less on higher management and more on their colleagues then the contributor may have concerns about appearing to be disloyal. In all of these cases, a natural reluctance can be compounded by a feeling of self-doubt. It may not be clear to the reporter that an adverse event has occurred. Those involved in an incident may seek to excuse or cover up their behaviour. Junior staff can also be reluctant to appear 'stupid' by raising concerns over unfamiliar equipment or procedures. As a result, they can remain silent about important safety concerns.

Many of the issues described above are illustrated by the events leading to the UK Bristol Royal Infirmary Inquiry. This focused on the procedures that were used to gain parental approval for child organ retention after autopsy. Concerns about these procedures were first identified following complaints that several complex cardiac surgical procedures continued to be conducted in spite of an unusually low recovery rate. The inquiry heard how Steve Bolsin, a member of staff within the unit, had attempted to draw attention to these problems by conducting a personal clinical audit. The following quotation comes from the hearings of this inquiry. The questions, labelled Q, were posed by the leagl team to the Chief Executive of the United Bristol Healthcare NHS Trust. His answers are labelled with an A.

> "Q. There was, was there, personal difficulty for a number of people in his overall conclusions being accepted?
>
> A. That certainly seems to be the case from all the records that I have seen, yes.
>
> Q. To what extent was that a reflection, would you say, of the absence of an institutionalised system of audit the absence of an institutionalised system of audit properly monitored, and to what extent did you consider that was part of a club culture where someone who rocked the boat, in whatever capacity, might be, as it were, going against the 'club'?
>
> A. They could both be contributory factors. Clearly, if there was no thorough-going structure in place along the lines we have discussed, then that is not going to lead to a climate whereby individuals doing audit and then presenting it is necessarily going to be received positively. Also, of course, if data is produced that appears to be critical of certain individuals and has not been collected with their knowledge and they do not subscribe to the methodology, then it would be surprising if they did not feel a degree of resentment and rejection of what was put in front of them. And it is possible that if this was undertaken by someone relatively new to the organisation who was challenging senior figures in the organisation, that, yes, indeed, it may have cut across some of the cultural boundaries within the Trust." [435]

In the subsequent investigations, Steve Bolsin's intervention was widely praised. However, things become more complex if an individual's actions can be interpreted as either 'whistler blowing' or 'trouble making' depending on ones' perspective. This dichotomy is illustrated by Mary Schiavo's criticisms of the FAA. She held the post of Inspector General in the US Department of Transportation. Following the Valujet crash, she told an American House of Representatives panel that she had made regular complaints to the FAA about what she felt were lax inspection practices in monitoring rapidly expanding airlines. Her comments and criticisms were widely reported in the media. However, her 'whistle blowing' was, in turn, heavily criticised by the US Congress. They attacked her by asking why she had not first passed her concerns to the Congress before publicly airing her criticisms. Under federal law, inspectors general are required to pass on to Congress within seven days any problems requiring immediate attention. She chose to resign from her post and subsequently published an account of her criticisms [729].

This dichotomy between constructive 'whistle blowing' and destructive criticism of an employer can also be seen in the Paul van Buitenen case. He voiced concerns about fraud and mismanagement in the European Commission's £60 billion budget. When these criticisms were made public, the veracity of his claims and his motivation for making them were, in turn, heavily criticised by individuals within the Commission. Although this incident did not have direct safety implications, his statements in a BBC interview provide a powerful illustration of the psychological pressures that affect such individuals:

> "I did not realise the full consequences of what would happen. I did not even know the word whistle-blower - I did not know this phenomenon existed... It was completely strange for me to see the commission tackle me on my personality and my credibility and not on the contents of what was disclosed. Sometimes I had difficulty keeping the tears inside when I discovered what machinery was brought against me... I am withdrawing as of April 1st, I want to be an anonymous official again. I want to show I can still be loyal, I want to do a normal standard budget management job. I want to have a quiet family life and be a husband and a father to my children who still have to do three years at secondary school, and I cannot carry on carrying this on my own." [103]

A UK National Audit Office enquiry headed by Sir John Bourn subsequently found errors totalling about £3 billion in European pay-outs during 1998. van Buitenen concerns are occasionally echoed in safety-related incident reporting systems: The provision of a reporting system is no guarantee of an appropriate safety culture in the companies that operate within an industry:

> "At the start of the Winter heavy maintenance programme, the company railroaded into place a computerised maintenance and integrated engineering and stores, planning and labour recording system. No training was given on the operational system only on a unit under test. Consequently we do not look at planes any more just VDU screens, filling in fault report forms, trying to order parts the system does not recognise, as the stores system was not programmed with (aircraft type) components (the company wanted to build a data base as equipment was needed)... The record had numerous faults, parts not recorded as being fitted, parts removed with no replacements, parts been fitted two or three times, parts removed by non-engineering staff, scheduled tasks not called-up by planning, incorrect trades doing scheduled tasks and certifying, and worst of all the record had been altered by none certifying staff after the CRS signatories had closed the work. Quality Airworthiness Department were advised of these deficiencies and shown actual examples. We were advised by the management that these problems are being addressed but they are not, we still have exactly the same problems today. What am I to do without losing my job and career. In a closed community like aviation, troublemakers and stirrers do not keep jobs and the word is spread around...' [174].

The comments that aviation is a "closed community" and that "troublemakers and stirrers do not keep jobs" provide an important 'reality-check' against some assertions about the benefits of incident reporting. These schemes have little impact on the underlying safety culture of many companies and organisations. O'Leary and Chappell argue that confidential incident reporting systems create a 'vital awareness of safety problems' [660]. The key point is not, perhaps, that O'Leary and Chappell are wrong but that the beneficial effects of these systems are constrained by the managerial culture in which they operate.

Media Disclosure

Issues of confidentiality and disclosure do not simply reflect the need to protect an individual's identity from their co-workers. They can also stem from concerns about media intrusion. For example, recent amendments have been proposed for ICAO Annex 13 on Accident and Incident Investigation and Prevention. The revisions would provide pilots with automatic confidentiality in accident and incident investigations. They would also limit the disclosure of information following an incident or accident. These amendments are significant in two ways. Firstly, they would ensure that the media had no right to cockpit voice recordings. This is an important issue given public and professional reactions to the broadcasting of such recordings after fatal accidents. Secondly, it would increase the level of civil protection available to pilots. The intention is to encourage a 'no-blame' approach to incident reporting. The concept is currently being tested in New Zealand civil courts.

If the ICAO adopts these amendments, it is likely that they will be ratified by all ICAO signatory nations as international law.

Accident and incident investigators often have a complex relationship with the media [419]. Public disclosure of sensitive information can jeopardise an enquiry and can dissuade contributions about potential hazards. Media interest can also play a powerful role in establishing reporting systems and in encouraging investment in safety initiatives. Peter Majgrd Nørbjerg's account of the new Danish Air Traffic Management reporting system reveals these two aspects of media involvement:

> "Then, in 2000, in order to push for a change the Chairman of the Danish Air Traffic Controllers Association decided to be entirely open about the then current obstacles against reporting. During an interview on national television, she described frankly how the then current system was discouraging controllers from reporting. The journalist interviewing the ATCO chairman had picked up observations made by safety researchers that, as described above, Denmark had a much smaller number of occurrence reports than neighbouring Sweden. Responding to the interviewer's query why this was so, the ATCO chairman proclaimed that separation losses between aircraft went unreported simply due to the fact that controllers - for good reasons - feared for retribution and disclosure. Moreover, she pointed out, flight safety was suffering as a consequence of this! These statements, broadcasted on a prime time news program, had the immediate effect that the Transportation Subcommittee of the Danish Parliament asked representatives from the Danish Air Traffic Controllers Association to explain their case to the Committee. Following this work, the Committee spent several of their 2000-01 sessions exploring various pieces of international legislation on reporting and investigation of aviation incidents and accidents. As a result of this, in 2001 the Danish government proposed a law that would make non-punitive, strictly confidential reporting possible." [676]

The irony in this account is obvious. The media played a key role in motivating political intervention to establish the reporting system. One of the first acts in establishing the new scheme was to create a legislative framework that effectively protected contributors from media exposure.

Proportionate Blame...

Potential contributors often have a justified fear of retribution. They may be dissuaded from participating in a reporting system if they feel that their colleagues and managers will perceive them to be 'whistle blowers'. Contributors can also be concerned about the legal consequences of submitting an incident report [83]. Leape points out that this reluctance is exacerbated by apparent inequities in the degree of blame that is associated with some adverse events. He also identifies a spectrum of blame that can lead from peer disapproval through to legal sanctions:

> "...these punishments are usually calibrated to the gravity of the injury, not the gravity of the error. The nurse who administers a tenfold overdose of morphine that is fatal will be severely punished, but the same dosing error with a harmless drug may barely be noted. For a severe injury, loss of the right to practise or a malpractice suit may result. Moderate injuries may result in a reprimand or some restriction in practice. Punishment for less serious infractions are more varied: retraining, reassignment, or sometimes just shunning or other subtle forms of disapproval." [480]

This fear of retribution has been addressed by number of regulatory organisations who have sought to ensure that any enforcement actions are guided by principles that are intended to protect individuals and companies. For example, the UK Health and Safety Executive is responsible for initiating prosecutions that relate to violations of health and safety law. These action are often taken in response to the accidents and injuries that are reported under the RIDDOR scheme, introduced in Chapter 1. The Health and Safety Commission requires that individual HSE inspectors inform their actions by the principle of proportionality; the enforcement action must reflect the degree of risk. They must also endeavour for consistency in their enforcement actions; they must adopt a similar approach in similar circumstances to achieve similar ends. A further HSE principle concerns the targeting of enforcement. Actions are focused on the people who are responsible for the risk and who are best placed to control it. Finally, there is a requirement that any legal or other enforcement actions should be transparent; the justifications and reasons for any decision to prosecute must be open to inspection. These guiding principles clearly distinguish regulatory actions from the informal retribution that often dissuades potential contributors from 'whistle-blowing'. In order to achieve these principles, Health and Safety inspectors will exploit a range of enforcement actions:

> "Enforcing authorities must seek to secure compliance with the law. Most of their dealings with those on whom the law places duties (employers, the self employed, employees and others) are informal - inspectors offer information, advice and support, both face to face and in writing. They may also use formal enforcement mechanisms, as set out in health and safety law, including improvement notices where a contravention needs to be remedied; prohibition notices where there is a risk of serious personal injury; withdrawal of approvals; variations of licences or conditions, or of exemptions; or ultimately prosecution. This statement applies to all dealings, formal or informal, between inspectors and duty holders - all contribute to securing compliance." [315]

The legal position of incident reporting systems is inevitably complicated by differences between different national systems. The effects of this can be seen from the differing reporting practices in European air traffic control. Some service provides are compelled to report all incidents to the national police force or to state prosecutors who will launch an investigation if they believe that an offence has been committed. However, there is a concern in the European coordinating organisation, EUROCONTROL, that controllers and pilots will significantly downgrade the severity of the incidents that they report in such potentially punitive environments. Concerns over litigation can also prevent reports from being filed. Other states have reached agreements between air traffic management organisations and state prosecutors to protect staff who actively participate in the investigation of an occurrence. The Swedish experience of operating an open reporting system is that very few controllers have lost their licenses as a result of filing an incident report within the last decade. The Luftfartsverket Flygtrafikjänsten personnel who operate the system stress the need to protect the controller's trust in the non-punitive nature of the system. The overall safety improvements from the information that is gathered by a non-punitive system are believed to outweigh the disciplinary impact of punitive sanctions. These arguments have also motivated the Danish system, mentioned earlier in this chapter [676]. It is interesting to note that the same personnel who expect a non-punitive approach to protect their submissions often also expect more punitive actions to be taken against others who are perceived to have made mistakes, especially pilots.

Most companies and regulators operate 'proportionate blame' systems. Annex 13 to the International Civil Aviation Organisation's International Standards and Recommended Practices provides the framework for accident and incident reporting in world aviation. This advocates a non-punitive approach to accident and incident reporting. It might, therefore, seem strange that some countries continue to operate systems that directly inform the actions of state prosecutors. There is, however, a tension between the desire to ensure the trust of potential contributors and the need to avoid a system that is somehow 'outside the law'. Ethical as well as judicial considerations clearly prevent any reporting system from being entirely non-punitive. For instance, action must be taken when reports describe drug or alcohol abuse. As a result most systems reserve the right to pass on information about criminal acts to the relevant authorities. This is illustrated by the immunity caveats that are published for NASA and the FAA's Aviation Safety Reporting System (ASRS). Section 5 covers the 'prohibition against the use of reports for enforcement purposes':

- "a. Section 91.25 of the Federal Aviation Regulations (FAR) (14 CFR 91.25) prohibits the use of any reports submitted to NASA under the ASRS (or information derived therefrom) in any disciplinary action, except information concerning criminal offences or accidents which are covered under paragraphs 7a(l) and 7a(2).

- b. When violation of the FAR comes to the attention of the FAA from a source other than a report filed with NASA under the ASRS, appropriate action will be taken. See paragraph 9.
- c. The NASA ASRS security system is designed and operated by NASA to ensure confidentiality and anonymity of the reporter and all other parties involved in a reported occurrence or incident. The FAA will not seek, and NASA will not release or make available to the FAA, any report filed with NASA under the ASRS or any other information that might reveal the identity of any party involved in an occurrence or incident reported under the ASRS. There has been no breach of confidentiality in more than 20 years of the ASRS under NASA management." [59]

Section 7 of the regulations governing the ASRS describes the procedure for processing incident reports. Again, this process explicitly describes the way in which legal issues are considered before reports are anonymized:

- a. "NASA procedures for processing Aviation Safety Reports ensure that the reports are initially screened for:
  1. Information concerning criminal offences, which will be referred promptly to the Department of Justice and the FAA;
  2. information concerning accidents, which will be referred promptly to the National Transportation Safety Board (NTSB) and the FAA; and Note: Reports discussing criminal activities or accidents are not de-identified prior to their referral to the agencies outlined above.
  3. time-critical information which, after de-identification, will be promptly referred to the FAA and other interested parties.
- b.Each Aviation Safety Report has a tear-off portion which contains the information that identifies the person submitting the report. This tear-off portion will be removed by NASA, time-stamped, and returned to the reporter as a receipt. This will provide the reporter with proof that he/she filed a report on a specific incident or occurrence. The identification strip section of the ASRS report form provides NASA program personnel with the means by which the reporter can be contacted in case additional information is sought in order to understand more completely the report's content. Except in the case of reports describing accidents or criminal activities, no copy of an ASRS form's identification strip is created or retained for ASRS files. Prompt return of identification strips is a primary element of the ASRS program's report de-identification process and ensures the reporter's anonymity." [59]

These quotations show that incident reporting systems must define their position with respect to the surrounding legislative and regulatory environment. They also illustrate the care that many organisations take to publish their position so that potential contributors understand the protection they are afforded. This does not necessarily imply that they respect the intention behind such protection. For instance, ASRS reporting forms are often colloquially referred to as 'get out of gaol free cards' by some US pilots.

The protection offered by confidential reporting systems has both positive and negative effects. 'No blame' reporting is intended to encourage participation in the system. Protection from prosecution can, however, introduce bias if it has greater value for particular contributors. This can be illustrated by the Heinrich ratios for US Commercial and General Aviation. The bottom tier of the Iceberg can be assessed through contributions to NASA's ASRS. Table 2.2 shows that General Aviation and air traffic management personnel submitted less voluntary incident reports than the crews of commercial air carriers in 1997 and 2000. These years were chosen because the ASRS provide complete month by month submission statistics. Administrative problems have led to submission data being merged for some months in other years. Others, including cabin crew, mechanics and military personnel provide very few submissions. The relatively high level of commercial aircrew contributions can be explained in terms of the protection offered by ASRS submissions. Submission to the system turns an adverse event into a learning opportunity In contrast, General Aviation pilots typically do not, typically, risk their livelihoods if their licences are revoked after an adverse event. There may also be less concern that others will witness and report an adverse event in General Aviation. They may, therefore, be less likely to submit information about adverse events they have been involved in. There is always the possibility in Commercial Aviation that other members of the flight crew or air traffic managers will file a report even if you do not.

	Air Carrier		General Aviation		Air Traffic Managers		Others	
	1997	2000	1997	2000	1997	2000	1997	2000
January	1,888	2,451	612	597	59	76	42	162
February	1,681	2,217	677	608	55	52	29	188
March	1,884	2,503	779	582	69	85	42	191
April	1,894	2,677	776	727	82	72	31	194
May	1,798	2,112	701	718	69	54	38	192
June	1,952	2,232	718	729	88	81	66	193
July	2,051	2,536	762	829	113	72	64	168
August	1,944	2,663	650	774	105	95	56	188
September	1,974	1,719	759	619	84	37	63	139
October	1,988	1,897	724	857	119	46	50	102
November	1,837	1,721	589	850	68	30	68	103
December	2,017	1,895	637	611	54	28	69	80
Total	22,908	26,623	8,384	8,501	965	728	618	1,900

Table 2.2: ASRS Contribution Rates 1997 and 2001

Table 2.3 presents NTSB data for accidents involving Commercial and General Aviation. In theory, this information can be used to calculate the Heinrich ratios that in turn illustrate the effects of ‘no blame’ reporting on participation rates. Unfortunately, the ASRS and NTSB use different classification schemes. The NTSB classify Commercial operations using the 14 CFR 121 and 14 CFR 135 regulations. In broad terms, 14 CFR 135 refers to aviation operations conducted by commuter airlines. 14 CFR 121 refers to larger air carriers and cargo handlers. The 14 CFR 135 statistics are further divided into scheduled and unscheduled services. Table 2.3, prsents the NTSB accident data for scheduled services. The 14 CFR 135 figures in parentheses also include accidents involving on-demand unscheduled services, such as air taxis. In calculating the Heinrich ratios, we have taken the figures for both scheduled and unscheduled services.

	14 CFR 135		14 CFR 121		General Aviation	
	All	Fatal	All	Fatal	All	Fatal
1997	16 (98)	5 (20)	49	4	1,845	350
1998	8 (85)	0 (17)	50	1	1,904	364
1999	13 (86)	5 (17)	51	2	1,906	340
2000	12 (92)	1 (23)	56	3	1,837	344
2001	7 (79)	2 (20)	45	6	1,726	325
2002	8 (66)	0 (17)	41	0	1,714	343

Table 2.3: NTSB Fatal and Non-Fatal Accident Totals

Figure 2.2 illustrates the Heinrich ratios for US Commercial and General Aviation in 1997 and 2000. The ratios were based on the number of incident submissions from Table 2.2. Table 2.3 provided the total number of fatal accidents. The number of non-fatal accidents was derived by subtracting the number of fatal incidents from the NTSB totals for all accidents. The General

Aviation classification is used in both the ASRS and NTSB statistical sources. The frequency of fatal commercial accidents was derived from the sum of incidents associated with 14 CFR 121 and 135 operations in the NTSB datasets. The figures in parentheses represent the total incident frequencies used in calculating the ratios.

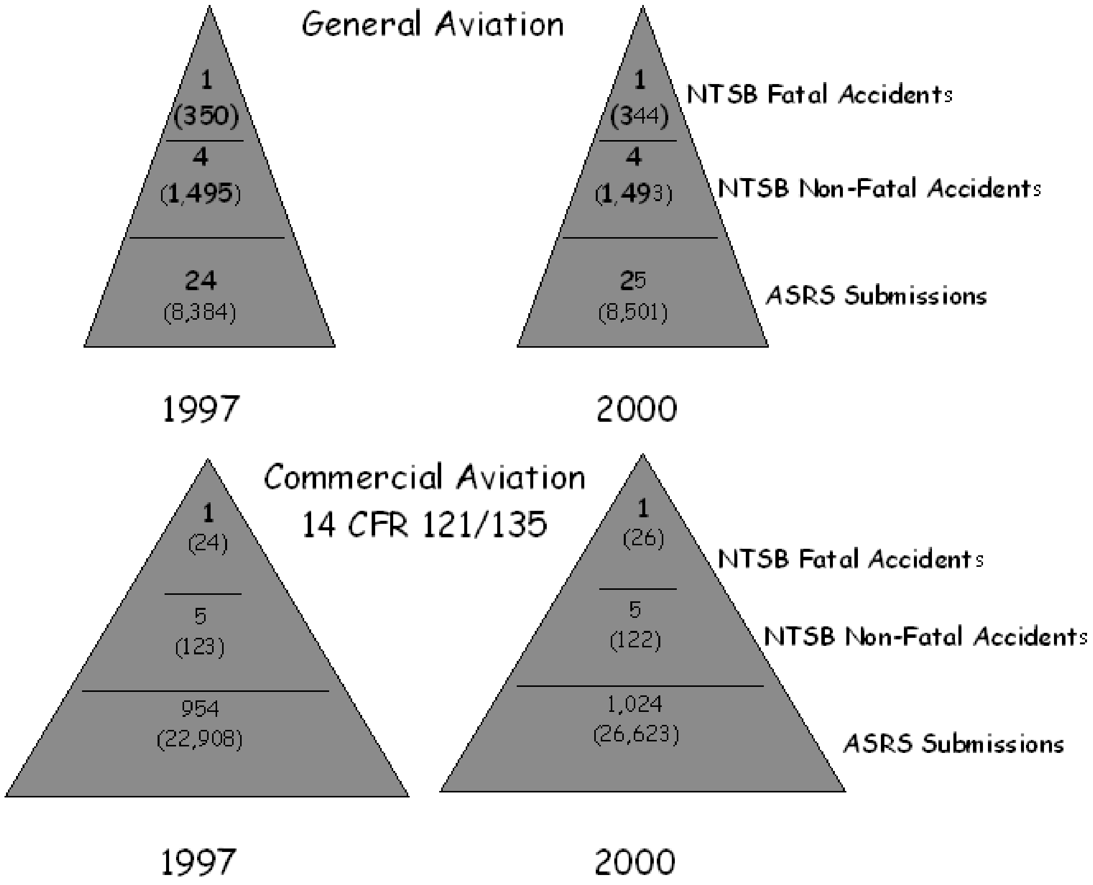

Figure 2.2: Heinrich Ratios for US Aviation (NTSB and ASRS Data)

The proportion of injuries to deaths in Figure 2.2 is lower for both General and Commercial Aviation than would be expected from Heinrich's ratio of one death to thirty injuries. In the case of General Aviation there is one fatal accident for every four non-fatal accidents. In Commercial Aviation, the ratio is one to five. This is deceptive. The ratios in Figure 2.2 cannot be directly compared to Heinrich's results. The NTSB and ASRS data refers to accidents rather than the number of injuries. The difference between Heinrich's ratio and our data arises because a single accident in the NTSB data can yield multiple fatalities or injuries. The NTSB do, however, present fatality and injury numbers for 14 CFR 121 operations. From this we can derive ratios of 1(2) : 10(21) : 13,311(26,623) in 1997 and 1(83) : 0.1(9) : 276(22,908) in 2000. The numbers in parentheses are the total frequencies for fatalities, minor injuries and incident reports. Further caveats can also be raised about these revised 14 CFR 121 ratios because the ASRS submission statistics combine 14 CFR 121 and 135 operations. These anomalies illustrate the practical difficulties that are often ignored by proponents of the Heinrich ratio as a tool for Safety Management. They also illustrate a recurrent observation in this book; incident and accident statistics are often presented in incompatible formats. This makes it difficult to trace the relative frequency of adverse events and their outcomes over time. It is apparent, however, that the revised 14 CFR 121 ratios are very different from Heinrich's figures. In particular the ratio of 1 death to 0.1 injuries seems at odds with the one to thirty ratio cited above. Fatal accidents are relatively rare in Commercial Aviation. Those that do occur often result in significant loss of life. Relatively few passengers and crew survive with minor injuries. These particular characteristics help to explain the apparent anomaly in the 1:0.1:276 ratio in 2000 for 14 CFR 121.

Heinrich's original work mixed outcome frequencies in terms of fatalities and injuries with event frequencies, based on observations of near misses. He did not attempt to estimate likely outcomes for near miss incidents. It can, therefore, be argued that the ratios in Figure 2.2 are more informative because they are based entirely on event frequencies. They do not include outcome information. Figure 2.2 can be used to identify patterns in ASRS submission data. In General Aviation, there was 1 fatal incident for every 24 submissions in 1997 and one fatal accidet for every 25 submissions in 2000. In Commercial Aviation, there were 954 ASRS submissions in 1997 and 1,024 in 2000 for each fatal incident. There are a number of possible explanations for these ratios. We can argue that there is a higher proportion of fatal accidents in General Aviation than in Commercial Aviation. This hypothesis is supported by the lower standards of training and equipment in General Aviation [82].

The higher rate of incident reports from Commercial Aviation in Figure 2.2 might be explained if these pilots had a greater incident exposure than in General Aviation. This is contradicted by the observation that General Aviation pilots accumulate significantly more flying hours than 14 CRF 121 and 14 CFR 135 operations combined. Table 2.4 presents NTSB statistics for flying hours and also accident rates per 100,000 hours in both Commercial and General Aviation [201]. To simplify the calculation of these rates we have excluded non-scheduled on-demand air taxis under 14 CFR 135. This is justified by the relatively low number of flying hours and incidents in this category.

	14 CFR 135		14 CFR 121		General Aviation	
	Accident Rate	Flying Hours	Accident Rate	Flying Hours	Accident Rate	Flying Hours
1997	1.628	982,764	0.309	15,838,109	7.19	25,591,000
1998	2.262	353,670	0.297	16,816,555	7.44	25,518,000
1999	3.793	342,731	0.291	17,555,208	6.4	29,713,000
2000	3.247	369,535	0.306	18,299,257	6.3	29,057,000
2001	2.330	300,432	0.231	17,752,447	6.28	27,451,000
2002	2.595	308,300	0.228	18,011,700	6.56	26,078,000

Table 2.4: NTSB Fatal and Non-Fatal Accident Rate Per 100,000 Flight Hours

The ratios in Figure 2.2 can also be explained in terms of a lower proportion of ASRS submissions from General Aviation than Commercial Aviation. Commercial pilots have more to lose from adverse events. The additional protection provided by the 'no blame' environment of the ASRS approach encourages them to submit a report. This attitude partly arises from the professional and personal consequences of losing a license that is essential to that person's job. Interviews with pilots have revealed that they are more likely to submit an ASRS report if they believed that someone else had also witnessed the incident. Given the NASA/FAA statement protection, cited above, there is perhaps a tendency to use the ASRS as a form of confessional in which contribution implies repentance. Arguably this has reached the point where many ASRS incidents are of a relatively trivial nature and provide few safety-related insights. With less to lose, General Aviation pilots may be less inclined to contribute to the system.

The difficulty in interpreting Heinrich ratios for US Commercial and General Aviation illustrates the confounding factors that must be considered when analysing reporting patterns. It seems likely that immunity policies affect contribution rates but little work has been conducted to determine how they interact with risk exposure, with individual attitudes to risk etc. The lack of such information is a primary motivation in writing this book. Major policy decisions have been made and continue to be made on the basis of data supplied by national and international reporting systems. There are, however, many open questions about the reliability, or biases, that affect these information sources.

2.3.2 Scope and Level

There are many different types of reporting system. Local schemes may record incident information supplied by a few staff in a particular department. International systems have been developed by groups such as the International Maritime Organisation to support the exchange of information between many different multinational companies [387]. These differences in the coverage of a reporting system can be explained in terms of their scope and level. The level of a reporting system is used to distinguish between local, national and international initiatives. The scope of a system defines the groups who are expected to participate in the scheme. The concept of coverage is a complex one. It is possible to distinguish between the theoretical and actual scope of a system. Although a system is intended to cover several different groups, such as medical and nursing staff, it may in practice only receive contributions from some subset of those groups. Similarly, a national system may be biased towards contributions from a particular geographical area.

There are important differences between national and regional reporting systems. For example, it can be easier to guarantee anonymity in national systems. Reports that are submitted to local systems often contain sufficient details for others to infer the identity of individuals who are involved in an adverse event. National systems are more likely to be protected by legal guarantees of confidentiality. They are also more likely to have the resources to finance technology protection for contributors, such as that offered by the Department of Energy's CAIRS system [655]. They can also finance dedicate personnel to process reports. Key individuals, such as the 'Gatekeepers' in the Swedish Air Traffic Control system, can be given the task of anonymizing information so that identities are hidden during any subsequent analysis. Steps may even be taken, as in the case of CIRAS, to ensure that these individuals are also prevented from retrieving identity information after the analysis is completed. All of these protection mechanisms are easier to sustain at a national level where resources of time, money and personnel can be deployed to address the logistical problems that often threaten locally-based systems.

A host of problems threaten anonymity in local reporting systems. For instance, the individuals who are responsible for setting up and running such a system can have some difficulty in convincing staff that they will not divulge confidential information to management or to other members of staff. One common means of avoiding this problem is to operate completely anonymous systems in which no identification information is requested. This creates the opportunity for malicious reporting in which one person implicates another. It also creates difficulties in both analysing and interpreting the causes and effects of particular incidents.

One of the longest running medical incident reporting systems was established in the Intensive Care Unit of an Edinburgh hospital. This scheme can be used to illustrate the difficulty of preserving anonymity and confidentiality in local reporting systems. The unit has eight beds [121]. There are approximately three medical staff, one consultant, and up to eight nurses per shift on the ward. Given the relatively close-knit working environment of an intensive care unit, it is possible for other members of staff to narrow down those individuals who might have submitted a report about a particular procedure or task that they were involved in. A key issue here is the trust that is placed in the person who is responsible for operating the system. The Edinburgh system was set up by David Wright, a consultant anaesthetist, who was heavily influenced by the earlier Australian Incident Monitoring Study (AIMS) [866]. This local system is heavily dependent upon his reputation and enthusiasm. He receives the reports and analyses them with the help of a senior nurse. The extent of his role is indicated by the fact that very few reports are submitted when he is not personally running the scheme.

The Paradox of Anonymity

There is a paradox in the affect that anonymity has on the value of a report at the local, national or international level. As part of the initiative to establish common guidelines for incident reporting in Air Traffic Control, interviews were conducted with controllers and other personnel in several European countries [423]. During these sessions, several contributors stressed the importance of anonymity. However, they also stressed the importance of knowing the context in which an incident occurred. This included both the location, which airport and which runway, as well as the time of day, the operator's shift pattern etc. Without this information, they argued that the report would have little or no value to other operators. With that information, however, it would be relatively easy to narrow the potential contributor down to a few individuals. The paradox here is that anonymity is often essential to encourage the continued submission of incident reports. However, anonymity jeopardises the usefulness of a report for those who may benefit most from the lessons that it contains.

In international schemes this paradox raises a number of deeper questions. A large number of local factors will influence the way in which an occurrence is dealt with. These include differences in national operating practices, in equipment, in workload. However, if a report were to be anonymized then much of this information would have to be omitted. It is not clear how much information about all of these issues ought to be provided and how much can be assumed about the readers knowledge of regional and national differences. In aviation, this has been addressed by ICAO Annex 13, mentioned above. This specifies the minimum content for accident and incident reports. However, these guidelines are not always adhered to. Similar provisions do not currently exist to support the sharing of data in the medical domain or in, for instance, rail transportation.

In local schemes, the context is already well established. The staff in the Edinburgh ICU system know that all reports refer to occurrences within that unit. As a result, much of the identifying information about that ICU can be reatined in the reports. Much of this detail would have to be removed in a confidential national systems in order to protect the individual hospital department. At the same time, however, there is an increased likelihood that those running local systems may be able to infer who contributed an anonymous report from their knowledge of the unit. The managers of the reporting system must ensure that similar inferences cannot easily be made by the co-workers who receive the recommendations that are generated from each contribution. This again leads to the danger that necessary information will be omitted.

'Targeting the Doable"

Local incident reporting systems must typically select their recommendations from a more limited set of remedial actions than national or international systems. For example, the FAA/NASA's ASRS is widely recognised to have a profound influence not just on US but also on global aviation policy. The same cannot be said for more local systems where it may only be possible to influence the unit in which it is being run. This is reflected in the more limited definition of an incident in some of these schemes. For example, the staff of the Yorkhill Hospital for Sick Children recently established an incident reporting system for incidents in a Neonatal Intensive Care Unit. This local system borrowed heavily from the existing schemes in Edinburgh and at various places in Australia [122, 121]. The agreed definition of an incident that fell within the scope of the system was printed on each of the forms:

> "A critical incident is an occurrence that might have led (or did lead) if not discovered in time - to an undesirable outcome. Certain requirements need to be fulfilled:
>
> 1. It was caused by an error made by a member of staff, or by a failure of equipment;
> 2. It can be described in detail by a person who was involved in or who observed the incident;
> 3. It occurred while the patient was under our care;
> 4. It was clearly preventable.
>
> Complications that occur despite normal management are not critical incidents. But if in doubt, fill in a form." [122]

The penultimate sentence illustrates a key point about local systems. Local schemes depend upon the good will, or at worst the passive acceptance, of higher levels of management. Such support can be jeopardised if the system is seen to move beyond constructive criticism.

Many of the incidents reported to local schemes can only be avoided or mitigated through cooperation with other, external organisations. For example, van Vuuren's study of incident reporting in a UK Accident and Emergency unit found that forty-five per cent of the causes (42 out of a total of

93) of the 19 incidents that were studies had organisational causes. Of these, thirteen causes were external to the Department itself. This is due to the way in which an Accident and Emergency department depends on the specialist services of other departments, including radiology, biochemistry etc:

> "Because the external factors are beyond the control of the investigated department, it is difficult to assess their real causes. It is of little use to hypothesise in detail about their origins and accompanying corrective actions of root causes in other departments... However, the majority of the external factors relate to the priorities of hospital management. The consequences of these priorities influence day to day practice in the A&E department, revolving mainly around staffing problems (not enough senior staff) and bed problems (lack of beds for A&E patients due to the continuous closing of beds on the wards), Although these external factors are beyond the control of the investigated department, their reporting is important to enable informal discussion between departments and to stimulate other departments to assess their own performance and its impact." [844]

There are clear differences between van Vuuren's emphasis on collecting data, even if it cannot immediately be used to affect other departments, and the previous definition of an incident which 'targets the doable'. The previous definition of a critical incident, arguably, illustrates the pragmatic approach that must be adopted during the establishment of an incident reporting system. Before the value of such a scheme has been widely accepted, it can provide difficult to get other groups to accept that their actions may lead to failures in the unit operating the system. Van Vuuren's argument that incident data can be used to enable informal discussions about common concerns will only be effective if other groups are willing to participate.

National and international systems can often make recommendations that have a much wider impact than local systems. For instance, the recommendations that are obtained from the UK's Royal College of Anaesthetists systems can be passed directly to other college's for further consideration[715]. Similarly, the GAIN system is intended to support the dissemination of 'best practice' across the World's airline operators and manufacturers [308]. It is also intended to support the dissemination of recommendations to air traffic service providers, airport managers etc. A number of limitations affect these large scale systems. It can be difficult to encourage the active participation of all regions within a system. These systems can also become victims of their own success if it becomes difficult to identify common patterns of failure amongst a large number of submissions.

Local, national and international systems provide different insights. For example, Section 2.1 described the potential benefits of incident reporting. These included the fact the they provide a reminder of hazards and that lessons can be shared. In a local system, these reminders may have greater local relevance than in a national scheme. In a national system, feedback often retains local features that were observed in the initial incident report. These features may not be appropriate for all participants. Alternatively, the incident must be abstracted to derive a generic account of the failure. In this case, the recipients must interpret the implications of the generic lesson in the context of their department or organisation. This can lead to a strongly negative reaction to the system if the lessons seem to be inappropriate [408]. There is also a danger of ambiguity; the implications of a generic lesson can be misinterpreted. The following list reviews a number of further differences:

- *local systems* can react relatively quickly to any report of an incident. As mentioned, the overheads of analysing and investigating a mishap can be substantially reduced because the individuals who run the system will have a good understanding of the context in which any failure occurred. These systems may only have a limited scope within a particular level of an organisation. Partly as a consequence of this, they often exploit ad hoc solutions to more serious problems. For instance, many hospital systems train their staff how to 'make do and mend' with poorly designed equipment [418]. National and international systems typically have the greater influence necessary to change procedures and prohibit the use of particular devices.

- *national systems* have correspondingly greater coverage. As a result, more reports may be received and better statistical data can be derived from them. This enables a closer relationship to be created between incident reporting and the subsequent risk assessments that drive future development and operational decision making. The ability to collate national data makes it more likely that such systems will be able to identify trends of common failures across many different sites. This is important because they can recognise the significance of what would otherwise appear as isolated failures. For instance, the lack of any effective central monitoring system has been identified as a reason why repeated problems with radiotherapy systems were not corrected sooner [487]. However, these systems introduce new problems of scale. There are considerable information processing challenges in identifying common trends in the 500,000 reports currently held by the ASRS . It can also be difficult to respond promptly when analysts must communicate with regional centres to establish the detailed causes of an adverse occurrence. Finally, it can be hard to ensure that local and regional agencies exploit consistent reporting procedures. This implies that similar incidents must be reported in a similar manner and that local or regional biases must be identified.

- *international systems* enable states to share information about relatively rare mishaps. They can also be used to exchange insights into the success or failure of recommendations for common problems. For example, Germany Air Traffic Control (Deutschen Flugsicherung GmbH) currently operates several parallel approach runways. The increasing use of these configurations has encouraged them to share data with other organisations which operate similar approaches, such as the UK's National Air Traffic Services operation at Heathrow. International reporting systems enable states to identify potential problems before they introduce systems that are currently operated in other countries. It can, however, be difficult to ensure the active participation of several different countries. Individual states must trust other countries both to investigate and report on their incidents. Cultural and organisational problems also affect the successful operation of international systems. For example, there is often a reluctance to adopt forms and procedures that were not developed within a national system. Occasionally, there is a belief that some of the incidents which are covered by national systems simply 'could not happen here' [423].

Large scale systems often attract political criticism if they are perceived to threaten other national and international organisations. It is for this reason that recent attempts to develop medical incident reporting systems in the United States are at pains to consider the relationship between federal and state bodies:

> "Congress should:
>
> - designate the Forum for Health Care Quality Measurement and Reporting as the entity responsible for promulgating and maintaining a core set of reporting standards to be used by states, including a nomenclature and taxonomy for reporting;
> - require all health care organisations to report standardised information on a defined list of adverse events;
> - provide funds and technical expertise for state governments to establish or adapt their current error reporting systems to collect the standardised information, analyse it and conduct follow-up action as needed with health care organisations.
>
> Should a state choose not to implement the mandatory reporting system, the Department of Health and Human Services should be designated as the responsible entity; and designate the Center for Patient Safety to:
>
> 1. convene states to share information and expertise, and to evaluate alternative approaches taken for implementing reporting programs, identify best practices for implementation and assess the impact of state programs; and
> 2. receive and analyse aggregate reports from States to identify persistent safety issues that require more intensive analysis and/or a broader-based response (e.g., designing prototype systems or requesting a response by agencies, manufacturers or others)."

[453]

The distinctions between local, national and international schemes often become blurred under systems such as that proposed by the US Institute of Medicine. Local initiatives report to State organisations that may then contribute to a Federal database. Such an integration will, however, change the nature of local systems. For instance, the need to ensure consistency in the information that is gathered nationally will force changes on the forms and procedures that are used locally. Recommendations that are issued from a national level may not easily be implemented under local conditions. For instance, recommendations relating to the use of more advanced equipment that has not yet been installed in all regions can serve to remind teams of what they are missing rather than forewarn them about the potential problems of equipment that they might receive in the future [409]. Similar comments can be made about initiatives to integrate national and international reporting systems [423]. The need to convert between national reporting formats and consistent international standards can lead to considerable tension. For instance, some European Air Traffic reporting systems operate a national system of severity assessment that must then be translated into categories proposed by EUROCONTROL's ESARR 2 document [717]. This translation process must be transparent if all of the member states are to trust the reliability of the statistics produced from international initiatives.

2.4 Summary

This chapter has summarised the reasons why a range of government and commercial organisations have established these systems in the military, in transportation, in healthcare, in power generation etc. These initiatives have been justified in terms of the learning opportunities that can be derived from incident data ideally before an accident takes place.

This chapter has also looked at some of the problems associated with incident reporting. These include the difficult of encouraging participation from a broad spectrum of contributors. For instance, we have calculated Heinrich ratios for fatal and minor accidents affecting US personnel. This reveals that contract staff may report fewer minor injuries than directly employed staff. The FRA have, therefore, encouraged greater monitoring of incidents involving contract workers.

'No blame' reporting systems encourage greater participation. However, the Heinrich ratios for General and Commercial Aviation suggest that the protection offered to contributors can introduce biases. In particular, pilots are more likely to report an adverse event if their livelihood is at risk or if they are concerned that their actions may be reported by colleagues and co-workers.

This book addresses the problems identified in this chapter. The aim is to present techniques that will help to realise the benefits that are claimed for incident reporting systems. Issues of anonymity, of legal disclosure, of retribution and blame, of scope and context must all be considered when developing an effective reporting scheme. It is also important to consider the sources of human error, system failure and managerial weakness that contribute to the incidents that are reported. This is the topic of the next chapter.

Chapter 3

Sources of Failure

Failures are typically triggered by catalytic events, such as operator error or hardware failure. These triggers exacerbate or stem from more latent problems, which are often the result of managerial and regulatory failure. In the most general sense, incident reporting systems provide a way of ensuring that such routine failures do not escalate in this manner. As a result, they must operate at several different levels in order to reduce the likelihood of latent failures and reduce the consequences of catalytic failures. Figure 3.1 provides an idealised view of this process. This diagram is a simplification. Political and economic necessity often break this chain of monitoring behaviour. Simple terms, such as "regulator" and "management" hide a multitude of roles and responsibilities that often conflict with a duty to report [773]. However, the following paragraphs use the elements of Figure 3.1 both to introduce the sources of failure and to explain why incident reporting systems

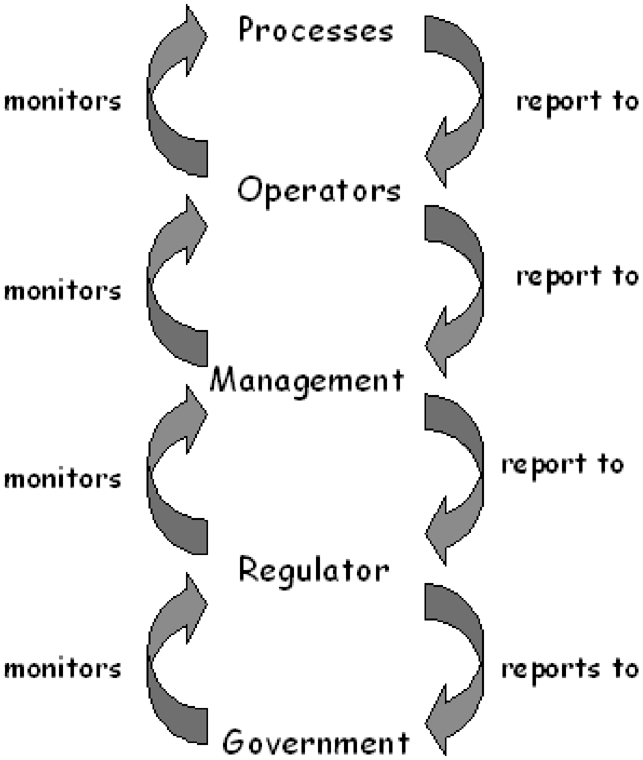

Figure 3.1: Levels of Reporting and Monitoring in Safety Critical Applications

have been introduced to identify and combat these sources once they have been identified.

3.1 Regulatory Failures

Regulation is centred around control of the market place. Regulators intervene to ensure that certain social objectives are not sacrificed in the pursuit of profit. These objectives include improvements in safety but they also include the protection of the environment, the preservation of consumer rights, the protection of competition in the face of monopolistic practices etc. For example, the Federal Railroad Administration's mission statement contains environmental and economic objectives as well as a concern for safety:

> "The Federal Railroad Administration promotes safe, environmentally sound, successful railroad transportation to meet current and future needs of all customers. We encourage policies and investment in infrastructure and technology to enable rail to realise its full potential." [239]

A similar spectrum of objectives is revealed in the Federal Aviation Administration's strategic plan for 2000-2001 [201]. The first of their three objectives relates to safety; they will 'by 2007, reduce U.S. aviation fatal accident rates by 80 percent from 1996 levels'. The second relates to security; to 'prevent security incidents in the aviation system'. The final aim is to improve system efficiency; to 'provide an aerospace transportation system that meets the needs of users and is efficient in the application of FAA and aerospace resources'.

3.1.1 Incident Reporting to Inform Regulatory Intervention

Regulatory authorities must satisfy a number of competing objectives. For example, it can be difficult to both promote business efficiency and ensure that an industry meets particular safety criteria. In such circumstances, regulatory duties are often distributed amongst a number of agencies. For example, the US National Transportation Safety Board (NTSB) has a duty to investigate accidents and incidents in road, rail and maritime transportation. All other regulatory activities in the field of aviation have been retained by the Federal Aviation Administration:

> "Congress (in enacting the Civil Aeronautics Act of 1938] is to provide for a Safety Board charged with the duty of investigating accidents... The Board is not permitted to exercise ... (other) regulatory or promotional functions. It will stand apart, to examine coldly and dispassionately, without embarrassment, fear, or favour, the results of the work of other people." (Edgar S. Gorrell, President, Air Transport Association, 1938 [482]).

The NTSB investigates the causes of incidents and accidents whilst the FAA is responsible for enforcing the recommendations that stem from these investigations. This separation of roles is repeated in other industries. For example, the US Chemical Safety and Hazard Investigation Board operates under the Clean Air Act. Section 112 (r) (6) (G) prohibits the use of the Board's conclusions, findings, or recommendations from being used in any lawsuit arising from an investigation. In contrast, the US Occupational Safety and Health Act of 1970 established the Occupational Safety and Health Administration (OSHA) to 'assure so far as possible every working man and woman in the Nation safe and healthful working conditions' through standards development, enforcement and compliance assistance.

Although the distinction between investigatory and enforcement functions is apparent in many different industries, the precise allocation of responsibilities differs greatly from country to country. For instance, the UK Rail Regulator is charged with safeguarding the passengers' interests within a 'deregulated and competitive' transportation system. However, the monitoring and enforcement of safety regulations remains the responsibility of the Railway Inspectorate. This differs from the US system in which the Federal Rail Administration takes a more pro-active role in launching safety initiatives. In the UK system, this role seems to rest more narrowly with the railways inspectorate that is directly comparable with the US NTSB.

We are interested in regulators for two reasons. Firstly, they are often responsible for setting up and maintaining the incident reporting systems that guide regulatory intervention. Secondly, regulators are ultimately responsible for many of the incidents that are reported by these systems. Similar failures that recur over time are not simply the responsibility of system operators or line managers, they also reflect a failure in the regulatory environment. Many regulators specifically have the task of ensuring that accidents and incidents do not recur. For instance, the US Chemical Safety and Hazard Investigation Board was deliberately created to respond to common incidents that were being addressed by 14 other Federal agencies, including the U.S. Environmental Protection Agency (EPA) and the U.S. Department of Labor's OSHA.

3.1.2 The Impact of Incidents on Regulatory Organisations

Regulators are increasingly being implicated in the causes of accidents and incidents [701]. In consequence, investigations often recommend changes in regulatory structure. The Cullen report into the Piper Alpha fire led to responsibility being moved from the Department of Energy's Safety Directorate to the Health and Safety Executive's Offshore Safety Division. Similarly, the Fennell report into the Kings Cross fire was critical of the close relationship that had grown up between the Railways Inspectorate and the London Underground management. Prior to Kings Cross, there had only been two Judicial Inquiries into UK railway accidents, the Tay Bridge disaster [357] and the Hixon Level Crossing Accident [549]. These criticisms reveal some of the problems that face regulators who must monitor and intervene in complex production processes. These problems can be summarised as a lack of information; a lack of trained personnel and a concern not to impose onerous constraints on trade.

Many industries increasingly depend upon complex, heterogeneous application processes. Most regulatory agencies cannot assess the safety of such systems without considerable help from external designers and operators. It is no longer possible for many inspectors to simply demand relevant safety information. They, typically, rely on the company and its sub-contractors to indicate which information is considered to be relevant to safety-critical processes. Rapid technical development, deliberate obfuscation, the use of (often proprietary) technical terms can all make it difficult for inspectors to gain a coherent view of the processes that they help to regulate. The activities of many regulatory agencies are further constrained by personnel limitations. These constraints partly stem from financial and budgetary requirements. It can be difficult to train and retain staff who are trained not only in the details of complex application processes but also in systems safety concepts. Even if it is possible to preserve a skilled core of regulators, it can be difficult to ensure that they continue to receive the 'up to date' training that is necessary in many industries.

Regulators must balance demands to improve the safety of complex application processes against the costs of implementing necessary changes within an industry. In 1999 Railtrack estimated that the cost of installing an Advanced Train Protection system over the UK rail network was in the region of £2 billion [690]. This system uses trackside transmitters to continuously monitor the activity of trains; including its speed, number of carriages, braking capacity etc. The ATP system will sense if the driver fails to react to any line-side instructions, including signals passed at danger, and will start to reduce the speed of the train. The costs of installing the more limited Train Protection Warning System was estimated by Railtrack to be in the order of £310 million. This system monitors the train before the key signals that protect junctions, single lines and 'unusual' train movements. A sensor is attached to the train and this detects emissions from two radio loops that are laid before these key signals. TPWS uses information about the current speed and the radio information that is transmitted when a signal is at red to detect whether the train is liable to stop in front of that signal. The information available to the system and the possible interventions are, therefore, more limited than ATP. The economic implications of regulatory intervention in favour of either ATP or TPWS are obvious. The Railway Safety Regulations (1999) require that ATP is fitted when 'reasonably practicable'. The wording of this regulation reflects the sensitivity that many regulators must feel towards the balance between safety and the promotion of commercial and consumer interests. If regulators were to recommend ATP rather than TPWS, rail operators would have been faced with significant overheads that many felt could not be justified by safety improvements. If they had recommended TPWS rather than ATP, passenger groups such as Railwatch would have criticised regulatory failure to introduce additional safeguards.

The Rand report was commission by the NTSB as part of an investigation into future policy for accident and incident investigation. This document questioned the nature of regulation in many safety-critical industries:

> "The NTSB relies on teamwork to resolve accidents, naming parties to participate in the investigation that include manufacturers; operators; and, by law, the Federal Aviation Administration (FAA). This collaborative arrangement works well under most circumstances, leveraging NTSB resources and providing critical information relevant to the safety-related purpose of the NTSB investigation. However, the reliability of the party process has always had the potential to be compromised by the fact that the parties most likely to be named to assist in the investigation are also likely to be named defendants in related civil litigation. This inherent conflict of interest may jeopardise, or be perceived to jeopardise, the integrity of the NTSB investigation. Concern about the party process has grown as the potential losses resulting from a major crash, in terms of both liability and corporate reputation, have escalated, along with the importance of NTSB findings to the litigation of air crash cases. While parties will continue to play an important role in any major accident investigation, the NTSB must augment the party process by tapping additional sources of outside expertise needed to resolve the complex circumstances of a major airplane crash. The NTSB own resources and facilities must also be enhanced if the agencys independence is to be assured." (Page xiv, [482])

A number of alternate models have been proposed. For instance, international panels can provide investgatory agencies with a source of independent advice. This approach is likely to be costly; such groups could only be convened in the aftermath of major accidents. In many industries, the dominance of large multi-national companies can make it difficult identify members who are suitably qualified and totally independent. Alternatively, investigatory agencies can develop specialist in-house investigation teams. The additional expense associated with this approach can make it difficult to also provide adequate coverage of the broad range of technical areas that must be considered in many incidents and accidents.

3.2 Managerial Failures

By failing to adequately address previous mishaps, regulators are often implicated in the causes of subsequent incidents. In consequence, they often help to establish reporting schemes as means of informing their intervention in particular markets. There are some similarities between regulatory intervention and the role of management in the operation of incident reporting systems. On the one hand, many organisations have set up incident reporting systems to identify potential weaknesses in production processes. On the other hand, many of the incidents that are reported by these schemes stem from managerial issues.

Social and managerial barriers can prevent corrective actions from being taken even if a reporting system identifies a potential hazard. These barriers stem from the culture within an organisation. For example, Westrum identifies a pathological culture that 'doesn't want to know' about safety related issues [862]. In such an environment, management will shirk any responsibility for safety issues. The contributors to a reporting system can be regarded as whistle blowers. Any failure to attain safety objectives is punished or concealed. In contrast, the bureaucratic culture listens to messengers but responsibility is compartmentalised so that any failures lead to local repairs. Safety improvements are not effectively communicated between groups within the same organisation. New ideas can be seen as problems. They may even be viewed as a threat by some people within the organisation. Finally, the generative culture actively looks for safety improvements. Messengers are trained and rewarded and responsibility for failure is shared at many different levels within the organisation. Any failures also lead to far-reaching reforms and new ideas are welcomed.

Westrum's categories of organisational culture mask the more complex reality of most commercial organisations. Accident and incident reports commonly reveal that elements of each of these stereotypes operate side by side within the same organisation. This is illustrated by the Australian Transport Safety Bureau's (ATSB) report into a fire on the Aurora Australis [48]. The immediate cause of the incident was a split fuel line to the main engine. Diesel came into contact with turbochargers that were hotter than the auto-ignition temperature of the fuel. It can be argued that the ship's operators resembled Westrum's bureaucratic organisation. Information about the modifications was not passed to the surveyors and other regulatory authorities. It can also be argued that this incident illustrates a pathological culture; ad hoc consultations perhaps typify organisations that are reluctant to take responsibility for safety concerns:

> "Consultations between the company and Lloyds Register and Wärtsilä, on the use of flexible hoses were ad hoc and no record of consultation or approval concerning their fitting was made by any party. No approval was sought from the Australian Maritime Safety Authority for the fitting of flexible hoses. Knowledge that the flexible hoses had been fitted under the floor plates was lost with the turn-over of engineers. The fact that other flexible hoses were fitted to the engines was well evident, but this did not alert either class or AMSA surveyors to the fact that the modifications were not approved." (Summary Conclusions, [48])

This same organisation also reveals generative behaviour. Persistent safety problems were recognised and addressed even if ultimately those innovations were unsuccessful. For instance, the operators of the Aurora Australis made numerous attempts to balance safety concerns about the fuel pipes against the operational requirements of the research vessel:

> "At an early stage of the ships life Wärtsilä Australia provided omega4 pipes to connect to the engines in an attempt to overcome the failures in the fuel oil pipework. This however did not solve the problem... When scientific research is being undertaken and dynamic positioning is in use, the isolation of noise and vibration from the hull is of importance. During these periods the main engines would not be in use. However the main generator sets are required and, to reduce vibration, the generator sets are flexibly mounted. For this reason, the generator sets were connected to the fuel system pipework with flexible hoses supplied by Wärtsilä. The subsequent approach in solving the problem on the main engines involved the fitting of sections of medium pressure hydraulic/pneumatic hose." (Page 33 - Engine Fuel Systems, [48])

Many investigators apply a form of hindsight bias when they criticise the organisational culture of those companies that suffer severe accidents. They have experienced a major failure and, therefore, these organisations must have a 'pathological' attitude to safety. This is over-simplistic. The previous incident has illustrated the complex way in which many organisations respond to safety concerns. It is possible to identify several different 'cultures' as individuals and groups address a range of problems that change over time.

3.2.1 The Role of Management in Latent and Catalytic Failures

MAnagement play an important role in the latent causes of incidents and accidents. The distinction between latent and catalytic factors forms part of a more general classification introduced by Hollnagel [362]. He identifies effects, or phenotypes, as the starting point for any incident investigation. They are what can be observed in a system and include human actions as well as system failures. In contrast, causes or genotypes represent the categories that have brought about these effects. Causes are harder to observe than effects. Their identification typically involves a process of interpretation and reasoning.

It is also useful to distinguish between proximal and distal causes [115]. In Hollnagel's terms, most incident reports focus on the proximal genotypes of failure. These the include 'person' and 'technology' related genotypes that are addressed later in this chapter. However, they also include 'organisation related genotypes' that address the role of line management in the conditions leading to an adverse event: "This classification group relates to the antecedents that have to do with the organisation in a large sense, such as safety climate, social climate, reporting procedures, lines of command and responsibility, quality control policy, etc." (Page 163, [362]). Hollnagel's classification of organisation genotypes reflects the increasing public and government interest in the distal causes of failure. He explicitly considers safety climate, social climate, *reporting procedures*, lines of command and responsibility and quality control policy as contributory factors in the events leading to failure. Table 3.2.1 illustrates how this high level categorisation can be refined into a check-list that might guide both the investigation of particular incidents and the development of future systems.

3.2.2 Safety Management Systems

Management can recruit a number of techniques to help them combat the latent causes of incidents and accidents. For example, Safety management systems help organisations to focus on "those elements, processes and interactions that facilitate risk control through the management process" [189]. The perceived success of this approach has led a number of regulators to support legislation that requires their use within certain industries, for example through the UK Offshore Installations (Safety Case) Regulations of 1992 . The UK Health and Safety Executive publish guidance material on the development of Safety Management Systems [319]. They emphasise a number of phases [189]:

- developing policy, which sets out the organisations general approach, goals and objectives towards safety issues;
- organising, which is the process of establishing the structures, responsibilities and relationships that shape the total working environment;
- planning, the organisational process which is used to determine the methods by which specific objectives should be set out and how resources are allocated;
- implementation which focuses on the practical management actions and the necessary employee behaviours that are required to achieve success;
- measuring performance, which incorporates the process of gathering the necessary information to assess progress towards safety goals; and
- auditing and reviewing performance, which is the review of all relevant information.

Incident reporting schemes offer a number of potential benefits within a safety management system. In particular, they can help to guide the allocation of finite resources to those areas of an application process that have proven to be most problematic in the past. In other words, incident reporting systems can focus risk assessment techniques using 'real world' reliability data that can be radically different from the results of manufacturer's bench tests. Incident reporting systems can also be used to assess the performance of safety management activities. They can provide quantative data that avoids subjective measures for nebulous concepts such as 'safety culture'. Managerial performance can be assessed not simply in terms of reduced frequency for particular incidents but also in terms of the reduced severity of incidents that are reported. Chapter 15 will, however, discuss the methodological problems that arise when deriving quantitative data from incident reporting systems.

3.3 Hardware Failures

Public attention is increasingly being focussed on the role of regulatory authorities in the aftermath of accidents and incidents. This has increased interest in incident reporting techniques as a means of informing regulatory intervention. Managerial failures also play an important role in creating the conditions that lead to many of the failures that are described in occurrence submissions. In consequence, a number of regulatory authorities have advocated the use of incident reporting techniques to help identify potential managerial problems within a wider safety management system. The following section builds on this analysis and begins to look at phenotypes and genotypes that relate to hardware failures. It can be argued that many of these failures stem from the distal causes of managerial failure. Stochastic failures can be predicted using probabilistic risk assessment. Design

General consequent	Specific consequent	Definition
Maintenance failure	Equipment not operational	Equipment (controls, resources) does not function or is not available due to missing or inappropriate management
	Indicators not working	Indications (lights, signals) do not work properly due to missing maintenance
Inadequate quality control	Inadequate procedures	Equipment/function is not available due to inadequate quality control
	Inadequate reserves	Lack of resources or supplies (e.g., inventory, back-up equipment etc.)
Management problem	Unclear Roles	People in the organisation are not clear about their roles and their duties
	Dilution of responsibility	There is not a clear distribution of responsibility; this is particularly important in abnormal situations.
	Unclear line of command	The line of command is not well defined and the control of the situation may be lost.
Design failure	Anthropometric mismatch	The working environment is inadequate and the cause is clearly a design failure.
	Inadequate Human-Machine Interface	The interface is inadequate and the cause is clearly a design failure.
Inadequate task allocation	Inadequate managerial rule	The organisation of work is deficient due to the lack of clear rules or principles
	Inadequate task planning	Task planning or scheduling is deficient
	Inadequate work procedure	Procedures for how work should be carried out are inadequate
Social pressure	Group think	The individual's situation understanding is guided or controlled by the group.

Table 3.1: Hollnagel's Categories for Organisational Genotypes

and requirements failures may be detected using appropriate validation techniques. However, many incidents defy this simplistic analysis of managerial genotypes as the root of all mishaps. Individual managers are subject to a range of economic, political and regulatory constraints that limit their opportunities to address potential hardware failures in many industries.

3.3.1 Acquisition and Maintenance Effects on Incident Reporting

Several factors affect the successful acquisition of hardware devices. Managers must have access to accurate reliability data. They must also be able to assess whether devices will be compatible with other process components. Compatibility can be assessed both in terms of device operating characteristics but also in terms of maintenance patterns. This is important if managers are to optimise inspection and replacement policies. A number of further characteristics must also be considered. The operating temperatures, humidity performance, vibration tolerances etc should exceed those of the chosen environment. components must meet electromagnetic interference requirements. They should also satisfy frequency, waveform and signal requirements as well as maximum applied electrical stresses. The tolerance drift over the intended life of the device should not jeopardise the required accuracy of the component. Finally, the component must fall within the allocated cost budget and must usually be available during the service life of an application process.

Many components fail to meet these requirements. Hardware failures have many different causes. The distal genotypes include design failures; the device may not perform the function that was intended by the designer. Hardware may also fail because of problems in requirements elicitation; the device may perform as intended but the designers' intentions were wrong. It can also fail because of implementation faults; the system design and requirements were correct but a component failed through manufacturing problems. A fault typically refers to lower-level component malfunction whilst failures, typically, affect more complex hardware devices. There are also more proximal genotypes of hardware failures. In particular, a device may be operated beyond its tolerances. Similarly, inadequate maintenance can lead to hardware failures. A number of military requirements documents and civilian standards have been devised to address these forms of failure, such as US MIL-HDBK-470A (Designing and developing maintainable products and systems) or the FAA's Code of Federal Regulations (CFR) Chapter I Part 43 on Maintenance, Preventive Maintenance, Rebuilding and Alteration. These standards advocate a number of activities that are intended to reduce the likelihood of hardware problems occurring or, if they do occur, to reduce the consequences of those failures. An important aspect of these activities is that they must continue to support the product throughout its operational life. Two key components of hardware acquisition and maintenance schemes are a preferred parts list and a Failure Reporting, Analysis and Corrective Action system (FRACAs). Preferred parts lists are intended to ensure that all components come from known or approved suppliers. These preferred parts lists also avoid the need for development and preparation of engineering justification for new parts and materials. They reduce the need for monitoring suppliers and inspecting/screening parts and materials. They can also avoid the acquisition of obsolete or sole-sourced parts. Failure Reporting, Analysis and Corrective Action systems provide individual organisations with a means of monitoring whether or not the components on a preferred parts list actually perform as expected when embedded within production processes in the eventual operating environment.

A continuing theme in this book will be that the use of safety-critical design and maintenance techniques, such as a preferred parts list, can have a profound impact on the practical issues involved in incident reporting. If a structured approach to hardware acquisition is not followed then it can be extremely difficult for engineers to effectively exploit the information that is submitted through a FRACA system. Engineers must assume that all components share similar failure modes even though they are manufactured by different suppliers. This can have considerable economic consequences if similar devices have different failure profiles, for example from different manufacturing conditions. Adequate devices may be continually replaced because of historic failure data that is based on similar but less reliable components. Conversely, it can be dangerous for engineers to assume that a failure stems from a particular supplier rather than from a wider class of similar devices. In order to support this inference, operators must analyse the different engineering justification for each of the different supplier's components to ensure that faults are not shared between similar devices from different manufacturers. The practical consequences of miscalculating such maintenance intervals is illustrated by work from the European insurance company Det Norske Veritas [458]. They assume that:

- that failure rate increases with increasing maintenance interval;
- that maintenance cost is inversely proportional to the maintenance interval
- that expected total cost is the sum of the maintenance cost and the expected failure cost.

It is possible to challenge these simplifying assumptions, however, they are based on considerable practical experience. Figure 3.2, therefore, illustrates the way in which the costs of maintenance are reduced as maintenance intervals are increased. It also shows the expectation that the costs of any failure will rise with increased maintenance intervals. The importance of this diagram for incident reporting is that each of these curves is based on the maintenance intervals and costs for particular devices. If a less reliable device were used with the same maintenance intervals then cost curves may be significantly higher, that is to say they will be translated along the Y-axis. Conversely, the cost curves for more reliable devices will be significantly lower even though the maintenance intervals will be based on less reliable devices. In either case, the effective use of reliability data for preventive maintenance depends upon the monitoring of devices from different suppliers within the actual operating environment of particular production processes [27].

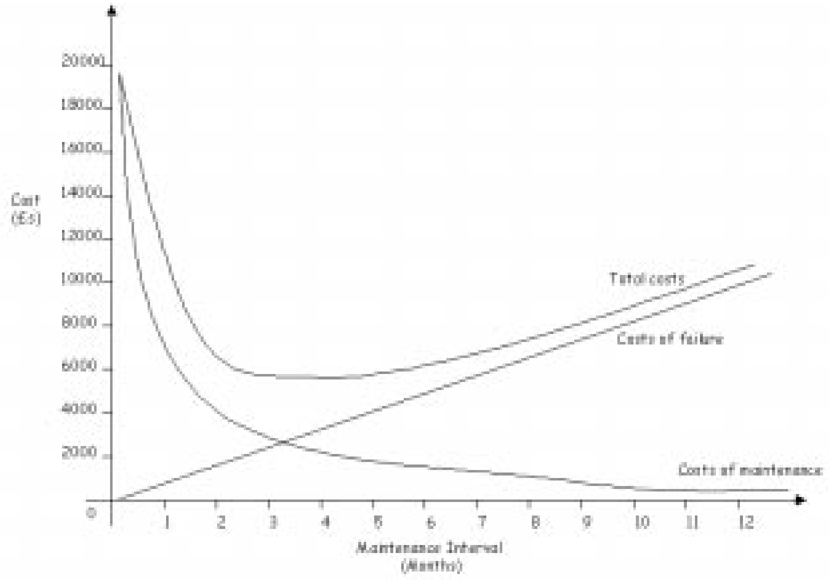

Figure 3.2: Costs Versus Maintenance Interval

3.3.2 Source, Duration and Extent

It is possible to identify a number of different types of hardware failure. In particular, they can be distinguished by their source, duration and extent [762]. Each of these failure types poses different challenges for the successful operation of incident reporting systems. The source of a failure refers to whether it is random or systematic. Component faults provide the primary cause of random hardware failures. All components have a finite chance of failing over a particular period of time. It is possible to build up statistical models that predict failure probabilities over the lifetime of similar devices. These probability distributions are usually depicted by the 'bath tub' curve shown in Figure 3.3. Initially there is an installation or 'burn-in' period when the component has a relatively

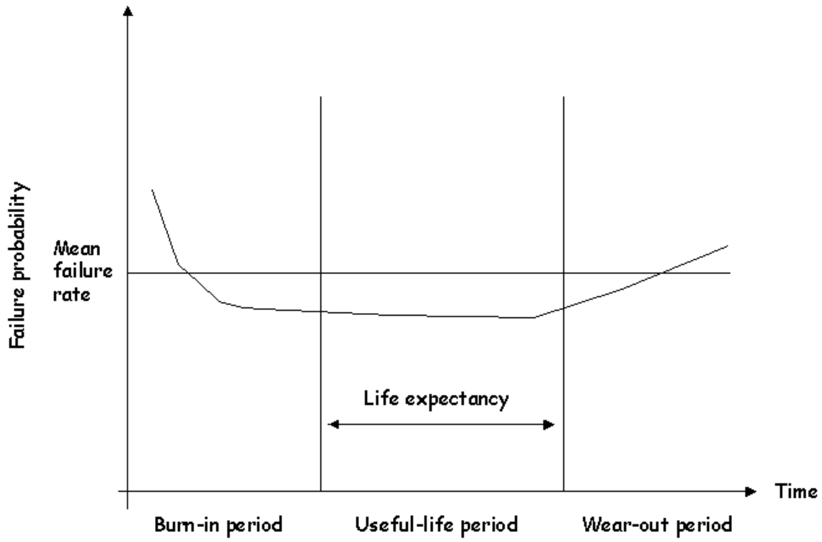

Figure 3.3: Failure Probability Distribution for Hardware Devices

high chance of failure. Over time, this declines for the useful life of the product until it begins to wear out. At this point, the likelihood of failure begins to increase. As can be seen from Figure 3.3 it is possible to abstract away from these lifecycle differences by suing a mean failure rate. However, this has profound practical consequences for the operation of an incident reporting system. When a class of components are first deployed, FRACAs submissions will indicate a higher than anticipated failure rate. This need not imply that the mean is incorrect, simply that the components must still go through the 'burning-in' period indicated in Figure 3.3.

The second source of hardware problems relates to systematic failures. These stem from errors in the specification of the system and from errors in the hardware design. Systematic failures are more difficult to combat using incident reporting techniques. The causes of particular mishaps may lie months or even years before a problem is reported by a supplier or end-user. It is for this reason that initiatives such as US MIL-STD-882D: Standard Practice for System Safety focus on the quality control and inspection procedures that are used throughout the design and implementation lifecycle. If systematic faults are found in hardware, or in any other aspect of a safety critical system, then this raises questions not just about the particular product that failed but also about every other product that was produced by that development process.

The duration of a failure can be classified as either permanent, transient or intermittent. Intermittent problems occur and then recur over time. For instance, a faulty connection between two circuits may lead to an intermittent failure. Occasionally the connection may operate as anticipated. At other times it will fail to deliver the correct signal. Conversely, transient failures occur once but may not recur. For instance, a car's starter motor may generate electromagnetic interference that will not recur until another car starts in the same location. Finally, permanent failures persist over time. Physical damage to a hardware unit, typically, results in a permanent failure. Each of these failure types poses different challenges for reporting systems. Transient failure can be particularly difficult to diagnose. They are, typically, reported as one-off incidents. This makes it very hard to reconstruct the operational and environmental factors that contributed to the failure. There is also a strong element of uncertainty in any response to a transient failure; it can often be very difficult for engineers to distinguish this class of failures from intermittent problems. The passage of time may convince engineers that a failure will not recur. This can be dangerous if the failure returns and proves to be intermittent rather than transient.

Permanent failures can seem simple to identify, diagnose and rectify. However, 'fail silent' components may leave few detectable traces of their failure until they are called upon to perform specific functions. Conversely, 'fail noisy' components may generate so many confounding signals that it can be difficult for engineers to determine which device has failed. It is important to stress that in practice there will seldom be a one-to-one mapping between each possible failure mode for any particular device and the reports that are submitted about those failures. For example, if two different members of staff identify the same failure then managers will be faced with the difficult task of working out whether or not those two reports actually do refer to the same problem or to two different instances of a similar failure. In such circumstances, it can take considerable time and resources for staff to accurately diagnose the underlying causes.

Intermittent failures are difficult to detect and resolve. Low frequency, intermittent failures may only be identified by comparing incident reporting systems from many different end-user organisations. The reports that document these failures may be distributed not only in time but also in geographical location. Many safety-critical products operate in similar environments in many different parts of the globe. Chapter 15 will argue that recent advances in probabilistic information retrieval and case based reasoning techniques for the first time provide effective tools for detecting and responding to this difficult class of failures. For now it is sufficient to observe that the identification of intermittent failures and trend information from incident reporting remains one of the biggest practical challenges to the effective use of these systems.

The final classification of failure types relates to the extent of its consequences. A localised fault may only effect a small sub-system. The consequences of a global fault can permeate throughout an entire system. Between these two extremes lie the majority of faults that may have effects that are initially localised but which, over time, will slowly spread throughout an application. In many instances it is possible to use incident reporting systems to chart the propagation of a failure over time. This provides valuable information not only about the failure itself but also about the reporting behaviour of the systems, teams and individuals who must monitor application processes.

The following incident report from the FDA's US Food and Drug Administration's Manufacturer and User Facility Device Experience Database (MAUDE) provides a glimpse of the complex relationships between device suppliers and the technical support staff who must operate them. In this case, end users made repeated attempts to fix problems that were created by the inadequate cooling of a patient monitor. The account of the problem clearly illustrates the end-user's sense of frustration both with the unreliability of the device and with the manufacturers' response:

> Monitors lose functions due to internal heat Note: several of the units returned for repair have had "fan upgrades to alleviate the temp problems". However, they have failed while in use again and been returned for repair, again salesman has stated it is not a thermal problem it is a problem with X's circuit board. Spoke with X engineer, she stated that device has always been hot inside, running about 68C and the X product has been rated at only 70C. Third device transponder started to burn sent for repair. Shortly after the monitor began resetting itself for no reason, fourth device monitor, SPO2 failed and factory repaired 10/01, 3/02. Also repaired broken wire inside unit 12/01. Tech 3/02 said the symptoms required factory repair... ([272], MDR TEXT KEY: 1370547)

This incident resulted in a series of follow-up reports. However, the manufacturers felt that the events described by the user could not be classified as safety-related; 'None of the complaints reported by the user were described as incidents or even near incidents. The recent report sent to the FDA appears to be related to frustration by the end user regarding the product reliability'. The manufacturer further responded by describing the evaluation and test procedures that had been used for each of the faulty units. The first had involved the customer replacing a circuit board. This did not fix the problem and the unit was sent back to the factory. The power supply was replaced but no temperature related failure was reproduced under testing by the manufacturers. A second device was also examined after a nurse had complained that the monitor had 'spontaneously' been reset. The hospital biomedical technicians and manufacturers representatives were unable to reproduce the transient failure and all functions were tested to conform to the manufacturers' specifications.

Manufacturers and suppliers are also often unable to determine the particular causes of reported mishaps. In the previous incident, the integrator/manufacturer believed that some of the problems might have stemmed from a printed circuit board made by another company. Tests determined that a board malfunction resulted in a failure to display patient pulse oxymetry waveforms on the monitoring system. The problems did not end when the integrator replaced the faulty board. The customer again returned the unit with further complaints that the device would not change monitoring modes. The integrator determined that the connectors to the printed circuit board were not properly seated. However, the board must have been properly placed prior to dispatch in order for the unit to pass its quality acceptance test. It is possible that the connector was not seated completely during the initial repair and gradually became loose over time. This incident illustrates the confusion that can arise when hardware devices are developed by groups of suppliers. The marketing of the device may be done by an equipment integrator who out-sources components to sub-contractors. For example, one company might provide the patient monitoring systems while another supplies network technology. This market structure offers considerable flexibility and cost savings during development and manufacture. However, problems arise when incidents stem from subcomponents that are not directly manufactured by the companies that integrate the product. Complaints and incident reports must be propagated back along the supply chain to the organisations that are responsible for particular sub-systems.

3.4 Software Failures

Software is now a key component in most safety critical systems. It is used to configure the displays that inform critical operating decisions, it can detect and intervene to mitigate the consequences of potential failures. Even if it is not used directly within the control loops of an application, it typically plays a key role in the design and development practices that help to produce the underlying systems. The Rand report into the investigatory practices of the NTSB emphasised the new challenges that these developments are creating:

> "As complexity grows, hidden design or equipment defects are problems of increasing concern. More and more, aircraft functions rely on software, and electronic systems are replacing many mechanical components. Accidents involving complex events multiply the number of potential failure scenarios and present investigators with new failure modes. The NTSB must be prepared to meet the challenges that the rapid growth in systems complexity posed by developing new investigative practices." [482]

The consequences of software-related incidents should not be underestimated. The failure of the London Ambulance Computer Aided Dispatch system is estimated to have cost between £1.1 and £1.5 million. Problems with the UK Taurus stock exchange program cost £75 to £300 million. The US CONFIRM system incurred losses in the region of $125 million [79] Few of these mishaps were entirely due to software failure. They were the result of "interactions of technical and cognitive/organizational factors than by technical factors alone" [533].

There are important differences between hardware and software failures. As we have seen, hardware failures can be represented as probability distributions that represent the likelihood of failure over the lifetime of a device. The practical difficulties of fabrication and installation prevent designers from introducing completely reliable hardware. If hardware related incidents exceed the frequency anticipated by the predicted failure probabilities then additional safeguards can be deployed to reduce the failure frequency or to mitigate the consequences of these failures. In contrast, software is deterministic. The same set of instructions should produce the same set of results each time they are executed. In consequence, if a software 'bug' is eliminated then it should never recur. There are some important caveats, however. In the real world, software operates on stochastic devices. In other words, subtle changes in the underlying hardware, including electromagentic interference, can cause the same set of instructions to have different results. In other applications, concurrent processors can appear to behave in a non-deterministic fashion as a result of subtle differences in the communications infrastructure [420]. Small differences in the mass of input provided by these systems may lead to radically different software behaviours. The problem is not that the code itself is non-deterministic. However, it can be almost impossible for operators and maintenance engineers to detect and diagnose the particular set of input conditions that caused the software to react in the manner that is described within an incident report. The consequences of this cannot easily be underestimated. In particular, it makes it difficult for engineers to distinguish between transient or intermittent hardware failures and software bugs arising from rare combinations of input conditions.

It can also be difficulty to ensure that bug fixes reach all end-users once a safety-critical product has been distributed. These practical difficulties are again illustrated by an incident report from the FDA's MAUDE system:

> "For approximately three weeks user hasn't been able to archive patient treatments due to software error. (The) facility has attempted to have company fix system in person but has only been successful at having company try by modem but to no avail." ([272], Report Number 269987)

The introduction of bug fixes can also introduce new faults that must, in turn, be rectified by further modification.

3.4.1 Failure Throughout the Lifecycle

Jeffcott and Johnson [396] argue that many software failures stem from decisions that are taken by high-level management. They illustrate this argument as part of a study into the organisational roots of software failures in the UK National Health Service. For example, the inquiry into the failure of the London Ambulance Computer Aided Dispatch System criticised the initial tendering process that was used:

> "Amongst the papers relating to the selection process there is no evidence of key questions being asked about why the Apricot bid, particularly the software cost, was substantially lower than other bidders. Neither is there evidence of serious investigation, other than the usual references, of Systems Options or any other of the potential suppliers' software development experience and abilities. ([772], page 18)

Such problems are typical of industries that are struggling to adapt management and procurement policies to the particular demands of software acquisition and development. They also illustrate the ways in which the various genotypes , such as managerial failure, help to create the conditions in which other forms of failure are more likely to manifest themselves.

The causes of software bugs can be traced back to the development stages where they were first introduced. For instance, the IEC 61508 development standard distinguishes between eleven lifecycle phases: initial conceptual design; the identification of the project scope; hazard & risk assessment; identification of overall safety requirements; resource allocation to meet safety requirements; planning of implementation and validation; system realization; installation and commissioning; validation; operation and maintenance; modification[420]. Software failures, typically, have their roots early in this development cycle. Many incidents stem from inadequate risk assessment. This is important in standards such as IEC 61508 that guide the allocation of software design resources in proportion to the predicted likelihood of a failure and its anticipated consequences. Errors during this risk assessment phase may result in unjustified attention being played to minor aspects of software functionality whilst too little care may be taken with other more critical aspects of a design. Any code that is then developed will fail to insure the overall safety of an application even though it runs in the manner anticipated by the programmer. Such problems are often caught during subsequent validation and verification. Those failures that do occur are, therefore, not only the result of an initial mistake or genotype. They also stem from failures in the multiple barriers that are intended to prevent faults from propagating into a final implementation. The IEC 61508 standard requires that the staff employed on each development task must be competent; they must understand the importance of their task within the overall development lifecycle; their work must be open to verification; it must be monitored by a safety management system; their ork must be well documented; it must be integrated within a functional safety assessment. These requirements apply across all of the lifecycle phases and are intended to ensure that failures do not propagate into a final implementation.

Managerial failures are an important precursor to other problems during software development, such as inadequate requirements capture [415]. This is significant because it has often been argued that the costs of fixing software bugs rise rapidly as development progresses. For example, Kotonya and Sommerville estimate that the costs of fixing a requirements error may be up to one hundred times the costs of fixing a simple programming error [459]. Such estimates have important implications for incident reporting. There can be insufficient resources to fix those software failures that are reported once a system is in operation. Many development organisations have introduced reporting schemes, such as NASA's Incidents, Surprises and Anomalies application, to elicit safety concerns well before software is deployed.

Requirements analysis helps to identify the functions that software should perform. It also helps to capture additional non-functional constraints; including usability and safety criteria. There are many reasons for the failure of requirements elicitation techniques. The following list provides a partial summary:

- *lack of stakeholder involvement.* The end-users who arguably know most about day to day operation may not be sufficiently consulted. In consequence, software engineers can get a distorted view of an application process. Similarly, some sectors of plant management and operation may not be adequately consulted. This may bias software engineers towards considering the requirements of one group of users' needs.

- *incorrect environmental assumptions.* A very common source of requirements problems stem from incorrect assumptions about the environment in which a software system will operate. Neumann's collection of computer related risks contains numerous examples of variables that have fallen above or below their anticipated ranges during 'normal' operation [627].

- *communications failures within development teams.* Incorrect assumptions about operating environments often occur because software engineers must often rely upon information provided by domain experts. Problems arise when these specialists must communicate technical expertise to people from other disciplines.

- *inadequate conflict management.* It is easy to underestimate the impact that social dynamics can have upon requirements engineering. Different stakeholders can hold radically different views about the purpose and priorities of application software. Requirements capture will fail if it does not address and resolve the tensions that are created by these conflicts. In particular, they can result in inconsistencies requirements, for example between speed and cost, that cannot be met by any potential design.

- *lack of 'ecological' validity.* It has increasingly been argued that requirements cannot simply be gathered by asking people about the intended role of software components [459]. in order to gain a deeper understanding of the way in which software must contribute to the overall operation of a system, it is important to carefully observe the day to day operation of that system.

As software engineering projects move from requirements elicitation towards installation and operation, they typically pass through a specification stage. This process identifies what a system must do in order to satisfy any requirements. It does not, however, consider the precise implementation details of how those requirements will be met. A similar array of problems affect this stage of software development:

- *inadequate resolution of ambiguity.* There is no general agreement about the best means of expressing requirements for large-scale software engineering projects. Formal and semi-formal notations provide means of reducing the ambiguity that can arise when natural language terms are used in a requirements document. However, these mathematical and diagrammatic techniques suffer from other limitations.

- *inadequate peer review.* Formal and semi-formal notations can be used to avoid the ambiguity and inconsistency of natural language. However, they may only be accessible to some of the people who are involved in the development process. In particular, they typically cannot be review by the domain experts and stakeholders who must inform requirements elicitation.

- *lack of change management.* Requirements will change over time as analysts consult more and more of the stakeholders involved in a system. These changes can result in 'feature accretion'; the core application functionality may become obscured by a lengthening wish-list of less critical features.
- *lack of requirements maintenance.* The constraints that software must satisfy will change during the lifetime of a system. Unless these changes trigger maintenance updates then software will continue to satisfy obsolete functional and non-functional requirements [434].

Errors in requirements elicitation and specification are more difficult to rectify than simple programming errors. There is, however, a bewildering array of potential pitfalls for the programmers of safety-critical systems. These include logical errors in calculations, such as attempting to divide a number by zero. They also include errors that relate to the handling of information within a program. For example, a variable may be used before it has been initialised with its intended value. The types of data that are represented within the program may not accurately match the full range of values that are provided as input to the program. The representations of these types may also differ between components of a program that are written by different teams or companies. The defences of strong typing that prevent such problems may be subverted or ignored. Valuable data may be over-written and then later accessed as though it still existed. A further class of problems relates to what is known as the flow of control. Instead of executing an intended sequence of instructions or of inspecting a particular memory location an arbitrary jump may be introduced through an incorrect reference or instruction. Other problems relate to the way in which a particular piece of code eventually executes at run-time. For example, there are differences between the precision with which data is represented on different target processors.

It is important not to underestimate the consequences of such coding errors. For example, the report into the London Ambulance Dispatch System failure records how such a bug caused the entire system to fail:

> "The Inquiry Team has concluded that the system crash was caused by a minor programming error. In carrying out some work on the system some three weeks previously the Systems Options programmer had inadvertently left in the system a piece of program code that caused a small amount of memory within the file server to be used up and not released every time a vehicle mobilisation was generated by the system. Over a three week period these activities had gradually used up all available memory thus causing the system to crash. This programming error should not have occurred and was caused by carelessness and lack of quality assurance of program code changes." ([772], page 45).

This quotation again illustrates the genotypes that lead to software failures. Errors can result from time and cost pressures; programmers may lack the necessary resources that are necessary to ensure type consistency and other necessary properties across module interfaces. If programmers receive inadequate training then they may fail to recognise that they have made an error. These problems can, in turn, be compounded by the lack of adequate tool support during various stages of implementation and testing.

Designers cannot be certain of eliminating all bugs from complex software systems. As a result, development resources must be allocated in proportion to the criticality of the code. If less resources are allocated to a module then there is, in theory, a higher likelihood that bugs will remain in that section of a program. Further problems stem from the difficulty of performing static and dynamic tests on complex and embedded systems. Dynamic testing involves the execution of code. This is intuitively appealing and can provide relatively direct results. It is also fraught with problems. It can be difficult to accurately simulate the environment that software will execute in. For instance, the Lyons report spends several pages considering the reasons why the inertial reference system (SRI) was not fully tested before Ariane flight 501:

> "When the project test philosophy was defined, the importance of having the SRI's in the loop was recognised and a decision was made (to incorporate them in the test). At a later stage of the programme (in 1992), this decision was changed. It was decided not to

> have the actual SRI's in the loop for the following reasons: the SRIs should be considered to be fully qualified at equipment level; the precision of the navigation software in the on-board computer depends critically on the precision of the SRI measurements. In the Functional Simulation Facility (ISF), this precision could not be achieved by electronics creating test signals; the simulation of failure modes is not possible with real equipment, but only with a model; the base period of the SRI is 1 millisecond whilst that of the simulation at the ISF is 6 milliseconds. This adds to the complexity of the interfacing electronics and may further reduce the precision of the simulation" (page 9, [505]).

Even in simple cases there are so many different execution paths and possible inputs that they cannot all be tested through dynamic analysis. As a result, many organisations have turned to combinations of both dynamic and static forms of testing. Static analysis evaluates the software without executing it. This relies upon reasoning about an abstraction of the specific machine that is eventually constructed by running code on a particular processor. For instance, walkthroughs can be performed by analysing the changing values of different variables as each line of code is executed by hand. Of course, this becomes increasingly problematic if the code is distributed. Formal, mathematical techniques can be used to reason about the behaviour of such software. However, all of these approaches rely upon reasoning about abstractions of the eventual system. There continue to be both theoretical and practical difficulties in refining proofs about models of a system into assertions about the potential behaviour of software operating on particular processors. The key point in all of this is that both static and dynamic testing provide means of increasing our assurance about the quality of a particular piece of code. Neither provide absolute guarantees. As a result, it seems likely that incident reporting systems will continue to provide valuable information about the symptoms of software failure for some time to come.

Redundancy can be used to reduce the likelihood of software failures. Several different routines can be used to perform the same function. The results from these computations can be compared and a vote taken to establish agreement before execution proceeds. If one section of code calculates an erroneous value then their result can be overruled by comparison with the other results. Lack of redundancy can, therefore, be seen to be a source of software failure. However, redundancy introduces complexity and can itself yield further implementation problems. It can also be difficult to ensure true diversity. For instance, programmers often resort to the same widely published solutions to common problems. If those solutions result in common problems then these may be propagated into several versions of the redundant code. Even if redundancy is successfully deployed, it can raise a number of further technical problems for the successful detection and resolution of incidents. For instance, redundancy is compromised if a routine continually computes an erroneous result but is successfully over-ruled by other implementations. The system will be vulnerable to failures in any of the alternative implementations of that function. It is, therefore, critical to monitor and respond to recurrent failures in redundant code.

Poor documentation can prevent technical staff from installing and configuring safety-critical applications. It can prevent end-users from responding appropriately to system prompts and directives. These problems can, in turn, compound the results of previous software failures if users cannot intervene in a timely fashion. Inadequate documentation can also be a cause of implementation errors in safety-critical programs. It is hard for programmers to correctly use their colleagues' work if they cannot understand the interfaces between modules. This problem also affects engineers who must maintain legacy systems. In particular, programmers often have to understand not simply what a piece of code does but also WHY it does it in a particular manner. This is critical if maintenance engineers are to justify their response to the problems identified by incident reporting systems. It is also important if engineers are to determine whether or not code can be deactivated or reused when it is ported between applications. There are close connections between these specific documentation issues, the problems of dynamic testing and the managerial causes of software failure:

> "Strong project management might also have minimised another difficulty experienced by the development. The developers, in their eagerness to please users, often put through software changes 'on the fly' thus circumventing the official Project Issue Report (PIR) procedures whereby all such changes should be controlled. These 'on the

> fly' changes also reduced the effectiveness of the testing procedures as previously tested software would be amended without the knowledge of the project group. Such changes could, and did, introduce further bugs." [772]

As mentioned, changes in the operating environment can invalidate the assumptions that were documented during any initial requirements engineering. Modifications that are introduced in response to those changes can, in turn, introduce further faults. Any one of these genotypes can lead to the incidents of software failure that are increasingly being documented by reporting systems[420].

3.4.2 Problems in Forensic Software Engineering

Many well-established techniques support the design and implementation of safety-critical systems. Unfortunately, very few support the investigation and analysis of software failure. These problems often manifest themselves in the recommendations that are made following such failures. In particular, many current standards advocate the importance of process measures as an indication of quality during safety-critical systems development. This means that regulators and quality assurance offices focus on whether appropriate practices have been followed during the various stages of the development process. They do not attempt to directly assess the quality of the final product itself. This avoids the many problems that arise when attempting to define appropriate measures of software quality [486]. However, this approach creates tremendous problems for the maintenance of incident reporting systems. The identification of a software fault throws doubt not only on the code that led to the failure but also on the entire development process that produced that code. At worst, all of the other code cut by that team or by any other teams practicing the same development techniques may be under suspicion. Readers can obtain a flavour of this in the closing pages of the Lyons report into the Ariane 5 failure. The developers must:

> "Review all flight software (including embedded software), and in particular: Identify all implicit assumptions made by the code and its justification documents on the values of quantities provided by the equipment. Check these assumptions against the restrictions on use of the equipment." [505]

Unfortunately, this citation does not identify any tools or techniques that might be used to 'identify all implicit assumptions' in thousands of lines of code. Such comments perhaps reveal some confusion about the practical problems involved in software development. This is illustrated by a citation from the report into the London Ambulance Computer Aided Dispatch system. Previous sections have identified a number of reasons why software cannot be totally reliable:

> "A critical system such as this, as pointed out earlier, amongst other prerequisites must have totally reliable software. This implies that quality assurance procedures must be formalised and extensive. Although Systems Options Ltd (SO) had a part-time QA resource it was clearly not fully effective and, more importantly, not independent. (Paragraph 3083, [772]).

Software-related incidents typically stem from more systemic problems. Bugs are often the result of inadequate funding or skill shortages. These failures are rooted in project management, including the risk assessment techniques that help to identify the criticality of particular sections of code. Many complex software failures also involve interactions between faulty and correct subsystems. They can stem from detailed interaction between hardware and software components. The nature of such incidents is illustrated by the following report from the FAA's Aviation Safety Reporting System. The erroneous TCAS II advisory interacted with the Ground Proximity Warning System:

> "Climbing through 1,200 feet [on departure] we had a TCAS II Resolution Advisory (RA) and a command to descend at maximum rate (1,500 to 2,000 feet per minute). [The flight crew followed the RA and began a descent.] At 500 feet AGL we leveled off, the TCAS II still saying to descend at maximum rate. With high terrain approaching, we started a maximum rate climb. TCAS II showed a Traffic Advisory (TA) without an altitude ahead of us, and an RA [at] plus 200 feet behind us... Had we followed the

> TCAS directions we would definitely have crashed. If the weather had been low IFR, I feel we would have crashed following the TCAS II directions. At one point we had TCAS II saying 'Descend Maximum Rate,' and the GPWS (Ground Proximity Warning System) saying 'Pull Up, Pull Up.' [The] ATC [Controller] said he showed no traffic conflict at any time." [546]

There are a number of reasons why traditional software engineering techniques cannot easily be applied to analyse the causes and consequences of software related failures. Most existing techniques address the problems of complexity by functional decomposition [486]. This assumes that by improving the reliability of individual components it is possible to improve the safety of an entire system. Such a decomposition often fails to account for interactions between subsystems. For example, the previous incident was caused by a software failure but resolved by operator intervention. Any redesign of the TCAS system must, therefore, ensure the reliability of the software and preserve the crews' ability to identify potential TCAS failures. A number of further problems complicate the use of traditional software engineering techniques to analyse incidents involving programmable systems. At one level, a failure can be caused because error-handling routines failed to deal with a particular condition. At another level, however, analysts might argue that the fault lay with the code that initially generated the exception. Both of these problems might, in turn, be associated with poor testing or flawed requirements capture. Questions can also be asked about the quality of training that programmers and designers receive. These different levels of causal analysis stretch back to operational management and to the contractors who develop and maintain application software. This multi-level analysis of the causes of software failure has important consequences. Existing software engineering techniques are heavily biased towards the requirements engineering, implementation and testing of safety-critical systems. There has been relatively little work into how different management practices contribute to, or compound, failures at more than one of these levels [396]. Leveson argues that:

> "...in general, it is a mistake to patch just one causal factor (such as the software) and assume that future accidents will be eliminated. Accidents are unlikely to occur in exactly the same way again. If we patch only the symptoms and ignore the deeper underlying cause of one accident, we are unlikely to have much effect on future accidents. The series of accidents involving the Therac-25 is a good example of exactly this problem: Fixing each individual software flaw as it was found did not solve the safety problems of the device" (page 551, [486]).

An alternative approach is to build on the way that standards, such as IEC61508, advocate the use of different techniques to address different development issues [879]. A range of different experts can be brought in to look at each different aspect of an incident. Management experts mght focus on the organisational causes of failure. Human factors specialists would use human factors techniques to investigate the role that operator behaviour played in an incident and so on. There are several objections to this approach. The cost of multidisciplinary investigations restrict them to high-risk mishaps. It can also be difficult to reconcile the views of individual team members from a range of different disciplines. Lekberg's has shown that the previous background of investigators will bias their interpretation of an incident [484]. Analysts are also most likely to finding the causal factors that are best identified using the tools and techniques that they are familiar with. In the case of software engineering, this might result in analysts identifying those causal factors that relate most strongly to requirements capture, to implementation or to testing rather than to the overall management of a software project. There is also a danger that such a multidisciplinary approach will suffer from problems that are similar to traditional techniques based on functional decomposition. If each expert focusses on their particular aspect of an incident then they may neglect the interactions between system components.

Further problems complicate the analysis of software failures. For example, simulation plays an important tool in many incident investigations. Several hypotheses about the sinking of the MV Estonia were dismissed through testing models in a specially adapted tank [227]. Unfortunately, incident investigators must often account for software behaviours in circumstances that cannot easily be recreated. The same physical laws that convinced the sub-contractors not to test the Ariane 5's inertial reference systems in the Functional Simulation Facility also frustrate attempts to simulate the incident [505]. Similarly, it can be difficult to recreate the exact circumstances which help to shape operator intervention. This is a general problem for the simulation of complex systems. However, it is particular severe for software systems that support synchronous interaction between teams of users and their highly distributed systems [415]. These issues form the focus of the next section.

3.5 Human Failures

Human failure plays a significant role in incidents and accidents. For instance, Van Cott cites studies which find that 85% of all incidents involving automobiles are caused by human error, 70% of all incidents in U.S. nuclear power plants, 65% in world wide jet cargo transport and 31% in petrochemical plants [185]. Similarly, Nagel argues that humans are implicated as 'causal factors' in more than half of all aircraft accidents. Within this figure, he argues they are involved in nine out of ten incidents involving general aviation [557]. These estimates can be misleading. Even those incidents that involve periodic hardware failures can be ascribed to human failure in the maintenance cycle. Failures that involve adverse meteorological conditions are caused by poor judgement in exposing the system to the risks associated with poor weather. It can be argued that all accidents and incidents are ultimately the responsibility of the regulatory authorities who must monitor and intervene to guarantee the safety of an industry. It is, therefore, perhaps better to distinguish between the proximal and distal impact of human error in the causation of adverse events. For instance, Heinrich claimed that up to 88% of all accidents stem from dangerous acts by individual workers [340].

3.5.1 Individual Characteristics and Performance Shaping Factors

Reason [699] and Wickens [863] provide sustained introductions to diverse forms of human error. In contrast, this section provides an introductory overview. ¡any reporting systems explicitly prompt investigators and respondents to identify what can be termed "performance shaping factors" [766] or the antecedents for error modes [362]. These factors can impair operator performance:

- *fatigue.* Incident reporting forms often ask specific questions about the shift patterns that operators and their colleagues worked immediate before the incident. Such information can be used to determine whether circadian rhythms, the natural variations in performance levels during the day, had any impact upon operator performance. For instance, Klein et al have shown that slight rhythmic variations can be seen in overall flying skills in each of the flight parameters over the time of day [447]. Worst performance was observed during the early morning. Hastings provides a review of more recent clinical work into the biological mechanisms that produce circadian rhythms [312]. He also provides a brief summary of the consequences that these mechanisms have for operator performance.

- *alcohol and drugs.* Tests for substance abuse are increasingly being conducted in the aftermath of incidents as well as accidents. Incident reports can also trigger increased workplace monitoring for drugs and alchol. This raises important ethical considerations for confidential systems. An increase in monitoring may compromise the identity of the individual or team who first raised concern about the issue. There are wider health and performance related issues. For example, it has been shown that short-haul aircrews significantly increase their alcohol consumption during periods away from home. This can increase heart rates during sleep which, in turn, has been shown to disturb the REM sleep that helps to determine sleep quality [293]. Caffeine and other stimulants are commonly used to compensate for the resultant fatigue.

- *stress.* Workplace stress stems from distractions, such as noise, but also to other environmental influences including heat, lighting levels as well as social pressures from colleagues. Sources of domestic stress include social pressures as well as financial and personal sources of anxiety.

Many studies have shown complex interactions between stress and performance. For instance, parachute jumpers have been shown to first improve their performance and then become worse at visual detection tasks as the time for their first jump approaches. It has also been shown that an individual's ability to detect changes in their environment becomes more focussed and that our ability to remember new information is impaired by increasing levels of stress [863].

- *workload.* Many reporting forms ask respondents to provide information about the number of tasks that operators had to perform immediately prior to an incident. They also often ask about differences in work patterns prior to an adverse event and about the division of responsibilities between members of a workgoup. All of these questions focus on the general mechanisms by which workload contributes to human error. Workload is, however, a nebulous concept. There are many different forms of measurement. Physical workload is relatively simple. It can be measured in terms of the oxygen consumption that operators require in order to convert the energy that is necessary to complete a given task [760]. Mental workload is more problematic. Wickens identifies a number of key questions about workload that can be adapted to guide incident investigation [863]. How busy was the operator? How complex were their individual or combined tasks? Is it reasonable to expect that additional tasks might have been handled above and beyond those already being performed? Did the operator respond to uncertain stimuli? How did the operator feel about the tasks being performed? Unfortunately, it can be hard to apply standard workload measures, such as NASA's Task-Load Index scale, in the aftermath of an incident [309]. Any subjective assessment of workload is likely to be influenced by the knowledge that a mishap has occurred.

- *individual differences.* Human resource managers have developed techniques to determine whether an individual is more or less likely to contribute to an accident. These tests examine character traits, including tendencies towards anxiety, fatigue, depression and boredom. They also consider age, gender, experience, personality traits and time sharing ability. One class of metrics considers what are termed 'learning styles'; these are important because there is no simple correlation between academic intelligence and ability in many diagnostic and control tasks [770]. Questionnaires have been developed to determine whether individuals are well suited to the acquisition and application of problem solving techniques. Such instruments can be applied post hoc, after an incident, to provide assurance that they are valid predictors of individual behaviour. However, this is arguably the most controversial form of measurement for any performance shaping factor or error inducing feature. The ethical implications are profound and problems of bias arise in the aftermath of an incident. In particular, it is difficult to separate individual differences as a cause of an incident from a myriad of other performance shaping factors. Incident information is not only used to validate personality questionnaires. It can also be used to drive simulations during training and selection exercises. For example, the FAA's Situation Assessment Through the Recreation of Incidents (SATORI) system is one of several that allows for the recreation of pre-recorded air traffic data through a controllers' plan view display and continuous readout update display for any sector [712]. This application was originally developed to recreate operational errors for review during quality assessment procedures but it has also been used to assess individual performance during the recreation of "error-inducing" situations.

- *attitudes towards risk.* We have defined risk to be the product of the probability of an incident and the seriousness of its consequences. The concept of risk is further complicated by uncertainty about the realisation of losses [506]. If an incident does occur then the actual consequences may depend upon a wide range of factors, including any mitigating actions taken by system operators. It is also possible to identify different individual attitudes towards risk taking that illustrate the underlying complexity of likelihood and consequence. For example, some individuals are risk averse whilst others actively seek exposure to certain hazards. Risk taking is the voluntary and conscious exposure to risk. Individual risk taking behaviour has often been cited as a factor behind the human contribution to incidents and accidents [722]. Higher speeds have been observed for drivers who have a previous record of accidents [856].

> Rockwell's pioneering study showed that electrical workers who take higher risks in their daily lives are also involved in more accidents at work [711]. There are, however, dissenting voices. Landeweerd et al have shown that the risk-taking tendency of construction workers was not related to a history of involvement in incident and accidents [474].

Hollnagel identifies many more of these performance shaping factors [362]. Their significance is that each factor can impair an individual's ability to call upon their perceptual, cognitive and physiological resources during the course of an adverse event. Physiology refers to the operator's physical attributes and includes their height, weight, reach etc. During an incident, operators can be temporarily incapacitated through injury or more permanently 'disabled' from performing their planned actions. Physiological failures can arise from barriers in the working environment; operators may not physically be able to reach a control. There are also more complex ways in which the body state of an operator can influence their performance. Teasdale and Barnard describe how physical conditions, such as heat or noise, can effect the mood of an operator. They go on to describe how such mood changes will also affect an individual's judgement [771]. Their work provides an analytical and theoretical explanation for the mass of empirical results that point to the increased likelihood of human error during operation in hot, noisy and cramped working environments [863]. Physiological problems directly lead to incidents if operators cannot complete planned actions. They may also indirectly lead to poor judgements and erroneous decisions through the cognitive mechanisms described by Teasdale and Barnard.

The majority of workplace accidents relate to collisions with moving and stationary objects. In 2000, the United Kingdom's Reporting of Injuries, Diseases and Dangerous Occurrences Regulations (RIDDOR) statistics record approximately 218 fatal injuried to workers [332]. Of these, falls from a height (29%), being struck by a moving vehicle (17%) or falling object (16%) are the most common form of injury. The non-fatal major injury rate for employees was approximately 120.1 per 100,000. Slips, trips or falls on the same level are expected to be the most common kind of non-fatal major injury to employees. The rate of injuried that resulted in an employee absence of 3 days or more was 21.4 per 100,000. Injuries sustained while handling, lifting or carrying are the most common kind of over-3-day injury to employees. There is a danger, however, that too much attention is paid to the immediate physiological impact of major incidents. Other long term physiological effects include functional aging. This is the deterioration of physical capacity beyond that which might be expected for the general population, that is to say beyond what might be expected from chronological aging. In particular, there is an increasing awareness that employers should also be concerned about the longer-term health and safety implications of particular tasks [863]. Many regulatory organisations are encouraging more active reporting of repetitive stress injuries, including carpal tunnel syndrome and work specific upper limb disorders. For instance, the US OSHA has proposed an ergonomic standard that is intended to prevent three million work-related musculoskeletal disorders over the next 10 years [651]. They estimate that such injuries currently cost $15 to $20 billion in workers' compensation costs with total costs as high as $45 to $60 billion each year. One of the key proposals in the OSHA standard is that companies should "set up a system for employees to report signs and symptoms of musculoskeletal disorders and respond promptly to reports".

Many physiological problems are caused by poor design [295, 157]. For example, Galer and Yap describe how existing input devices make data entry errors more likely in patient monitoring systems within an intensive care unit [282]. Junge and Giacomi describe how some of these problems have been addressed during the development of the general purpose workstation on the space shuttle [433]. Many physiological problems stem from operator behaviour. Workers in many industries, including car production, marine engineering and electricity generation, neglect risk reducing measures [311]. They ignore many of the dangers associated with incorrect postures or with unbalanced positions. Risk-taking is viewed as a controllable part of their everyday life at work [368]. There are other sources of physiological injury within the workforce. Studies of incidents involving postal workers have also shown that supervisors may expose their colleague to situations, such as adverse weather conditions, that significantly increase the risks of an injury[77]. Incident reporting systems, such as that proposed by the OSHA standard, have been advocated as a means of addressing these problems. Newsletters can disseminate information about previous mishaps that involve the violation of guidelines on appropriate posture . Direct information about real incidents often proves to be more effective than abstract classroom-based training sessions [425].

Perceptual failures occur when operators fail to correctly detect important cues and signals from the environment. For instance, the crews' apparent inability to sample critical information from an engine vibration montor was identified as a causal factor in the Kegworth crash [8]. Many authors have commented on the clutter that characterises the cockpits of modern commercial aircraft [863]. Billings notes, however, that there is a tension between filtering information to reduce the perceptual loading on operators and actively hiding information that may be essential for fault diagnosis [82]. The specific problems of cockpit design are also reflected in other industries. Sheridan describes a loss of coolant incident in a nuclear reactor that caused more than five hundred annunciators to change status in the first minute and more than eight hundred within the first two minutes [738]. In contrast, some safety-critical systems provide operators with too little information about the state of an application. Cook and Wood cite a medical incident report to illustrate this potential cause of human 'failure':

> "During a coronary bypass graft procedure, an infusion controller device delivered a large volume of a potent drug to the patient at a time when no drug should have been flowing. Five of these microprocessor-based devices were set up in the usual fashion at the beginning of the day, prior to the beginning of the case. The initial sequence of events associated with the case was unremarkable. Elevated systolic blood pressure (> 160 torr) at the time of the sternotomy prompted the practitioner to begin an infusion of sodium nitroprusside via one of the devices. After this device was started at a drop rate of 10/min, the device began to sound an alarm. The tube connecting the device to the patient was checked and a stopcock (valve) was found to be closed. The operator opened the stopcock and restarted the device. Shortly after the restart, the device alarmed again. The blood pressure was falling by this time, and the operator turned the device off. Over a short period of time, hypertension gave way to hypotension (systolic pressure <60 torr). The hypotension was unresponsive to fluid change but did respond to repeated boluses of neosynephrine and epinephrine. The patient was placed on bypass rapidly. Later the container of nitroprusside was found to be empty; a full bag of 50mg in 250ml was set up before the case". [182]

An experienced physicians had set up this device so that it allowed a free flow of the drug into the patient once the physical barrier of the stopcock was removed. A visual and an auditory alarm were presented when the device was started because there was no flow with the stopcock closed. When the stopcock was opened, the same alarms were again presented. This time, however, the device could not detect drops being administered because the drug was passing freely into the patient. Their blood pressure dropped and so the physician shut-down the device. However, this did not prevent the continued flow of the drug. Such incidents emphasise that we cannot isolate our ability to perceive an alarm from our ability to detect the additional information that is necessary to diagnose the causes of the alarm. In the reactor's loss of coolant incident the operator was overwhelmed by the sheer number of information sources, in the medical mis-administration incident they failed to detect any information that might have helped form a more correct diagnosis of the problem.

		State of the World	
		Signal	Noise
Response	Yes	Hit	False Alarm
	No	Miss	Correct Rejection

Table 3.2: Outcomes from Signal Detection

Environmental factors affect our ability to perceive information. High ambient noise levels can prevent operator from hearing particular warnings. On the other hand, attempts to overcome ambient noise levels have led some developers to produce warnings that reach up to 100 decibels at the pilot's ear. Such sound levels are likely to have a profound impact upon an individual's ability to attend to, or process, other information [668]. Some sources of environmental interference are less easy to predict than high ambient noise levels:

> "[On takeoff], at approximately 500 feet AGL, a laser beam of green light struck through the right side window of my cockpit striking my First Officer in the right eye and blinding both he and I for approximately 510 seconds due to the intensity of the light beam. I immediately notified the Tower Controller [who stated] that this had become a recurring problem with the laser show coming from the top of the [hotel] in Las Vegas. We were very fortunate, because this could have been a much more serious situation had the laser struck myself as well as [my First Officer] at a more direct angle, severely blinding both of us and endangering the lives of my passengers and crew." [667]

Many mishaps stem from problems of signal detection. Table 3.2 explains some of the issues involved in this aspect of perception. If a signal is present then either the operator may detect it, in which case they have achieved a 'hit', or they may fail to detect the signal, this results in a 'miss' in Table 3.2. If the signal is absent and the user detects it then this results in a false alarm. There are also situations, especially in the medical domain, where it may be better for the patient to act as though a signal were present even though there may be some uncertainty about the observation [281]. The final alternative is that the operator correctly observes an absent signal.

Other forms of perceptual failure arise from the difficulty of correctly sampling many different items of information. This is not simply a problem in using foveal and peripheral vision to scan a large number of displays, it also relates to the rate at which information changes over time. De Keyser has studied the impact of these temporal issues in domains ranging from steel production to healthcare [437, 438]. Operators are liable to miss critical information if it is rapidly replaced by other signals. Conversely, they are unlikely to detect trends that emerging over hours, days or weeks, especially if their attention is diverted by other tasks. This is typified by incidents of involving navigational failures. An initially small degree of error gradually grows with potentially disastrous consequences, as in this grounding reported by the Australian Maritime Incident Investigation Unit. The Pilot's likelihood of detecting the error was decreased by the fact that he was presumed to be asleep during part of the passage:

> "The ship continued on a gyro heading of 354 degrees to make good a course of 350 degrees at a speed of about 13.8 knots. The state of tide was about two hours before low water and what tidal stream there was tended to set the ship to the east. The 2nd mate fixed the ships position at 02:49 and again at 03:07, when about 3 nm from Heath Reef. Both positions put the ship to the east of the intended course line. The weather was fine with some cloud, the wind was from the south-east at 18 - 20 knots. There was only one vessel, a fishing vessel, in the vicinity of Heath Reef, which was showing a broad red side light. At about 0311, the 2nd mate touched the pilot on the shoulder to remind him to make the scheduled mandatory report to Reef Centre. The pilot got down from the chair and picked up the VHF radio and duly reported the ships position and speed. As he looked forward at Heath Reef, he realised that New Reach was in the wrong relative position. He ordered an alteration of course to 350 degrees. The pilot could also see the fishing vessel, but it was well clear of New Reach. However, the skipper of the fishing vessel used channel 16 VHF to contact New Reach and inquired whether the pilot wanted him to pass New Reach to starboard (green to green). The pilot replied that it was not necessary and that he was just dodging around Heath Reef..." [521]

An operator's ability to sample information can depend upon the mode of presentation. There are some obvious differences. For example, auditory displays typically have a shorter temporal duration than visual displays. Conversely, it can be easier to filter individual sounds from a large number of simultaneous auditory signals than it is to detect individual changes in a bank of visual displays. There are also a number of less obvious properties. For instance, Posner, Nissen and Klein point to the dominance of auditor warnings over visual alarms [685]. Both audio and proprioceptive, or tactile, alarms provoke faster responses than visual warnings. However, operators more reliably provide the response associated with the visual alarm if they are faced with both an auditory and a visual warning. If an auditory task is being performed concurrently with a visual one then the auditory task tends to suffer most from this division of attention.

Wickens provides an excellent overview of the ways in which the human perceptual system contributes to, and helps to avoid, major incidents [863]. Perception cannot, however, be isolated from other attributes of human behaviour. The way in which an individual will attend to different information sources is heavily determined by cognitive or mental models of the processes being observed. If operators think that a process is about to enter a hazardous state then they will, typically, devote additional perceptual resources to monitor Norman [636] argues that the development of appropriate models can be supported by the provision of appropriate feedback about system behaviour. that process. Unfortunately, these mental models are not always accurate. Operators often fail to predict critical changes in an application process. There is a lag between any increase or decrease in process error rates and any appreciable change in human sampling. Sheridan builds on this analysis [737]. He argues that the time between two observations of an instrument should be determined by a cost-benefit trade-off between growing uncertainty about the state of an unsampled channel and the costs of sampling that channel. The main practical problem with this analysis is that both of these estimates are likely to be highly subjective. For example, an expert may be able to predict the state of a process variable with far greater certainty than a novice. A risk adverse individual may also associate greater costs with NOT sampling a channel than a risk preferring individual.

Cognition refers to the ways in which we process the information that we perceive in our environment. An operator's perception of a signal or warning is influenced by their mental model of an application. Cognition and perception are, therefore, closely inter-twined. This is illustrated by the following NTSB incident report in which an AMTRAK express collided with a Maryland commuter train. The engineers believed that a the signal 1124-2 was on CLEAR when it was actually set to APPROACH. This persuaded him not to pay special attention to the subsequent signal at Georgetown junction. His mental model of the state of the track made him anticipate a clear line and this directed his perception of critical indications to the contrary:

> "The APPROACH indication of signal 1124-2 required the MARC train 286 engineer to slow his train to not more than 30 mph after passing the signal and to be prepared to stop at the Georgetown Junction signal. The collision occurred because the engineer did not operate MARC train 286 in conformity with the signal indication when he stopped at Kensington station and then proceeded towards Georgetown Junction, attaining a speed of about 66 mph. The engineers actions after departing the Kensington station were appropriate had signal 1124-2 been CLEAR, but his actions were inappropriate for an APPROACH aspect...
>
> If the engineer thought that his last signal (1124-2) was CLEAR, none of the signals he could have normally expected at Georgetown Junction would have been so restrictive as to demand his immediate action. Hence, he had no reason to try to see the signal as soon as possible. In addition, there was no radio conversation between train engineers and the dispatcher that could have provided the MARC train 286 engineer with a clue on the other trains operating in the area. Disbelief was likely once he or the other crewmembers or both observed the STOP signal at Georgetown Junction. The crew would have then consumed some time trying to reconcile the restrictive STOP indication with an expected CLEAR indication, which had been the norm for them at Georgetown Junction. One of the passengers stated, I could see the look, like bend over and check to see if somethings coming, then they jump back like in shock, then they went forward again just to double check, which would attest to disbelief on the part of the traincrew." [596]

This incident clearly indicates the strong connections between cognition, in terms of memory and use of mental models to inform expectation, and perception, in terms of sampling critical information. Teasdale and Barnard extend this analysis to show further interaction between physiology and both cognition and perception [771]. The physical 'well being' of an operator not only affects their ability to perceive critical information, it can also prevent them from acting effectively on that information, for example in situatiuons of extreme cold or noise. Figure 3.4 provides a high level overview of the way in which cognition can affect these diverse aspects of human behaviour. As mentioned before, the perception of information about the current state of the system can be biased by our prior beliefs about what are, and what are not, salient sources of information that must be sampled. Our analysis of the information that we perceive can also be biased. For example, there is a strong tendeny to recognise information that confirms previous expectations and to ignore contradictory indications. Kletz describes an example of this form of bias:

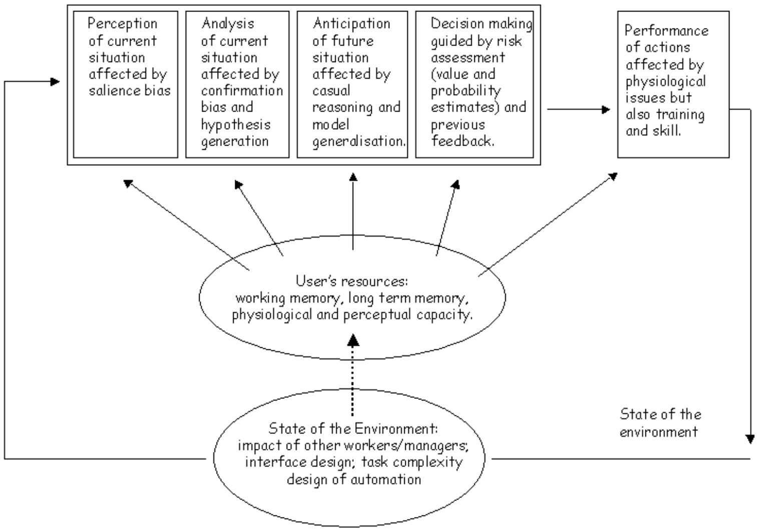
Perception of current situation affected by salience bias
Analysis of current situation affected by confirmation bias and hypothesis generation
Anticipation of future situation affected by casual reasoning and model generdlisation.
Decision making guided by risk assessment (value and probability estimates) and previous feedback.
Performance of actions affected by physiological issues but also training and skill.
User's resources: working memory, long term memory, physiological and perceptual capacity.
State of the Environment: impact of other workers/managers; interface design; task complexity design of automation
State of the environment

Figure 3.4: Cognitive Influences in Decision Making and Control

> "The operator correctly diagnosed that the rise in pressure in the reactor was due to a failure of the ethylene oxide to react. he decided that the temperature indicator might be reading high and that the temperature was, therefore, too low for reaction to start or that the reaction for some reason was sluggish to start and required a little more heat. he, therefore, raised the setting on the temperature trip and allowed the temperature to rise. (Two people were injured by the resulting explosion). His diagnosis, though wrong, was not absurd. However, having made a diagnosis he developed a mind-set. That is, he stuck to it even though further evidence did not support it. The temperature rose but the pressure did not fall (the reaction was exothermic). Instead of lloking for another explanation or stopping the addition of ethylene oxide, he raised the temperature further and continued to do so until it reached 200 degrees C instead of the usual 120 degrees C." [449].

There are several different forms of confirmation bias. For example, many people seem to exploit a representativeness heuristic. This favours familliar hyptheses that match the set of symptoms which we observe in our environment. Problems arise when the symptoms are similar to, but not an exact match, for those typically associated with a hypothesis. Under such circumstances, there is a tendency to select the familliar hypotheses rather than considering the probability of competing diagnoses [863]. Similarly, the availability heuristic describes how some hypotheses are more easily brought to mind than others. For instance, Javaux's work on pilot interactions with flight management systems has identified both recency and frequency effects that biasindexbias!frequency bias their expectations about the modes that are exhibited by these applications [394]. Fontenelle argues that incidents which are described in greater detail to the workers in safety-critical applications will also be perceived as having a greater prior probability [251].

Figure 3.4 also shows how an operator's analysis of the current situation is affected by their anticipation of future states. Such predictions are based on mental models that reflect our understanding of application processes. Such an understanding will always be simplistic and incomplete for all but the most rudimentary of systems. The following case from the Swiss Critical Incidents in Anaesthesiology system illustrates how correct mental models not only depend on an understanding of the basic functionality of a system, but also on the particular characteristics of system design. An incomplete understanding of the oxygen flush on an inhalator led to incorrect predictions about the induction of an anaesthesia:

> "During induction of inhalational anaesthesia (50% N2O/50% O2/sevoflurane up to 8 Vol%) the patient did not reach a sufficient level of anaesthesia (there was only a superficial anaesthetic level with profound agitation which could be achieved although a sevoflurane oncentration up to 8 Vol% was used). The anaesthetic machine (Carba) was tested in the morning by the nurse and was found to be working correctly. During the event, the oximeter showed a FiO2 of near 75%, although a fresh gas mixture of 2 l N2O/min and 2 lO2/min. was choosen and could be seeen on the rotameters. Surprisingly, the ventilation bag of the circle-circuit didn't collapse during inspiration and the boy didn't pass the excitation phase of the induction. A anaesthetic gas analyzer was not used. Because there must have been a surplus of fresh gas, the machine was checked again and the problem was found: this type of old anaesthetic machine has a oxygen flush button, which MUST TURNED ON AND MUST BE TURNED OFF AFTER USE. So, during checking the machine in the morning, the O2-flush button was tested, but not completely turned off again, so that the bypassed oxygen diluted the sevoflurane and the fresh gas mixture. Correcting this problem, the anaesthetic was completed successfully and with no further problem. The saturation of the patient was never below 97%." [755]

Figure 3.4 illustrates how decision making is linked to the operator's perception of the current situation, to their analysis of that situation and to predictions about the potential future situation. Such decision making is determined by implicit assumptions both about the benefits of particular actions and the likelihood of obtaining those benefits. The resulting decisions cannot simply be characterised in terms of numerical comparisons between the products of these two terms. Individual attitudes to risk and the perception of potential benefits can lead to a number of well known paradoxes that are confirmed by incident reports:

> "Suppose a physician sees 48 breast cancer patients per year. Two treatments are possible, with the following outcomes predicted: if treatment A is prescribed, 12 patients will survive. If treatment B is prescribed, there is a 0.25 probability that 48 patients will survive and a 0.75 probability that no patients will be saved. Which treatment would you prescribe if you were a physician? although, the estimated outcome is identical most people given such a choice choose treatment A, the sure thing, over B the calculated risk. " [446]

Figure 3.4 is intended to show that an operator will iterate between the stages involved in perceiving the current situation, analysing that situation, predicting future situations and eventually making a decision. It also illustrates how the operator's mental and physical resources can have a profound impact upon their ability to perform each of the phases described in previous paragraphs. For example, fatigue might impair an operator's ability to accurately perceive necessary signals in their environment. Similarly, high demands on working memory might lead them to form an incorrect assessment of their current situation even though they may have identified necessary information. These cognitive, perceptual and physiological resources are, in turn, affected by the operator's environment. Noise, heat, vibration can have physiological impacts upon a worker. The inefficient allocation of tasks, poor interface design or interruptions from colleagues can stretch cognitive and perceptual resources. Some of these factors act directly on the feedback loop between the operator's actions and their perception of the environment in Figure 3.4. Other factors such as managerial or domestic pressures may act to influence operator behaviour in a less direct manner. This is denoted by the dotted line from environmental influeces to the elipse representing operator's resources in the diagram.

3.5.2 Slips, Lapses and Mistakes

Errors can be seen as the unwitting deviation of actions from intentions. Operators may forget to perform a necessary command or they may repeat unnecessary steps. Errors can also be seen as the unwitting deviation of planned actions from a goal. Operators may mistakenly believe that certain actions will lead to a desired outcome. This definition of error ignores the important question of goal formation. It does not describe the many complex ways in which training, the presentation of display information, intervention from colleagues or other factors in the working environment help to shape the strategies and objectives that determine our more immediate objectives. For instance, Gaba has outlined a number of ways in which anticipation helps to shape strategy formation and goal setting [281]. He then uses this analysis to describe the knock-on effects that can emerge when inappropriate strategies help to 'provoke' the more detailed forms of error referred to in the previous definitions. Hollnagel also describes how human reliability will decline as operators move from strategic and tactical modes of control to opportunistic and scrambled interventions [362]. These different control modes have a strong impact upon intentions and actions that lead to errors.

Errors do not occur in a social or regulatory vacuum. They occur against a background of rules, regulations and procedures. Violations, therefore, are the deliberate contravention of those practices that are necessary to preserve the safety of a system. From this it follows that an error need not be a violation and that a violation need not involve an error. The violation of an inappropriate rule may be necessary to preserve the safety of an application process. Duncan describes an incident in the North Anna reactor that illustrates such a necessary violation [219]. Changes in the generation process led to dangerous temperature profiles following a scram. This process involves the insertion of neutron absorbing control rods to reduce reactivity. The operators were faced with a difficult choice. Following the Three Mile Island accident, Nuclear Regulatory Commission (NRC) regulations required that operators delay any intervention in order to allow a more detailed situation assessment during any potential emergency. However, plant management believed that if they obeyed this regulations then the safety of the plant would be threatened. They would no longer be able to predict its behaviour. If they disobeyed the regulations then the plant could be saved but they would beak the NRC conditions of operation. The plant management chose to violate the regulation; a pump was taken off the coolant circuit and the emergency was resolved. Duncan observes that this incident underlines the dangers of trying "to prescribe regulations, procedures or algorithms, especially when these prescriptions are backed by legal sanctions" [219].

If an individual does not know that they are violating a rule or procedure then this can be interpreted as an error. Unfortunately, incident investigators cannot always discern the intentions of an operator. If intention can be demonstrated then it is possible to identify three different types of deliberate violation. The North Anna example, cited above, illustrates the more general class of necessary violations [701]. Such incidents illustrate situations in which rules and regulations place staff in danger. In contrast, a routine or normal violation is one which involves some element of 'corner cutting'. This is typical of situations in which a group of skilled worked accept dangerous working practices as the norm. A good example, would be the removal of necessary protection devices. Finally, an optimising deviation involves some form of personal gratification or thrill seeking. An individual may deliberately choose to ignore accepted operating practices in order to 'optimise the joy of speed or indulge in aggressive instincts' [701].

Many incidents stem from complex combinations of optimising, necessary and routine violations. This is illustrated by a US Chemical Safety and Hazard Identification Board investigation of an incident at an explosive company:

> "The investigation team found that operators regularly used metal tools to unplug mixing pot draw-off lines in Booster Room 1. Several explosives manufacturing incidents during melt/pour operations at other companies have been caused by using metal tools to chip or forcefully break apart clogs in draw-off valves... The plant manager found (one of these tools) in Booster Room 1 on more than one occasion. When the manager found the rod in the booster room, he stated that he told operators not to use the tool, and the rod was taken to the tool room. Operators reported, however, that this tool was routinely kept in Booster Room 1 and was also used to push unmelted TNT on the

> surface down into the liquefied TNT in the melting pots. Operators indicated that it was sometimes very difficult to clear valves, so they had to use more force. The metal rod would be jammed into the valve repeatedly until the mass of material was broken free. The tool would have to be extracted quickly when the clog was freed because the hot, melted explosive mixture would flow from the open valve stem and would burn the worker clearing the valve if the worker was not fast enough. Being burned by the molten liquid was considered to be the primary hazard associated with this activity." [160]

From the perspective of the manager, the use of the tool was a routine or normal violation. In contrast, the workers may have viewed the same violation as a necessary means of completing their tasks on schedule and without exposing themselves to what they perceived to be the primary hazard. The workers' justification for violating the managers instructions was based upon a mistaken judgement about the primary hazard. The consequences of an explosion were greater than being burned by the molten liquid. This example also illustrates the problems of investigating the violations that contribute to incidents and accidents. Violations are strongly connected to ideas about operating norms. The use of the metal tools was 'normal' practice within the work group. It was an abnormal violation for the management and regulators. From this it follows that any member of the work group who reports on this 'normal' violation will be seen as a whistle-blower or someone who violates the norms of their working group. Chapter 5 describes a number of techniques that can be used to overcome the natural reticence of workers to report on the potentially dangerous working practices of their colleagues. Some managers collude in optimising violations. In such circumstances, the reporter or whistle blower must not only be assured of their anonymity but also of the independence of any subsequent investigation.

In the same way that we can distinguish between necessary, optimising and routine violations, it is also possible to identify different types of errors. The most general classifications separate slips and lapses from mistakes. Slips involve failures in the execution of a plan. They often have observable consequences, such as a slip of the tongue. Lapses involves failures in a well understood sequence of actions regardless of whether that plan was appropriate. They describe more covert forms of error, including failures of memory such as forgetting someones name. They may only be apparent to the person experiencing them. Slips and lapses can be distinguished from mistakes. Mistakes are failures of intention rather than execution. They stem from a failure to select appropriate objectives irrespective of whether the actions taken to achieve those objectives were successful.

Reason's Generic Error Modelling (GEMS)[699] brings together the slip, lapse and mistake taxonomy with Rasmussen's Skill, Knowledge and Rules approach to cognition [694]. Skill-based performance takes place after the statement of an intention or objective and is characterised by a lack of conscious control. It is typical of expert interaction, is smooth and appears to be automated. Knowledge and rule based performance only occur after an operator is made aware of a potential problem. Rule based performance occurs when individuals meet familiar problems that can be resolved through the recall and application of rules and procedures. Knowledge based performance typifies interaction in unfamiliar situations where operators must consciously rely upon inference and stored knowledge to identify a solution. Slips and lapses mainly occur during skill based performance. Inadvertent errors of omission or commission are likely during the unconscious pursuit of a recognised objective [363]. In contrast, errors of rule based performance are liable to result in mistakes. For instance, an operator may not identify the problem at hand and, therefore, select rules and procedures that are more appropriate to other problems. Alternatively, if a user correctly identifies the state of the system then they may apply the wrong rules and procedures. Users either apply bad rules or misapply good rules. Errors at the knowledge based level are also likely to result in mistakes. For example, operators may pursue inappropriate objectives if they possess incomplete, inconsistent or incorrect knowledge about their system. This can be caused by thematic vagabonding in which operators flit from one aspect of a problem to another without pausing to conduct a sustained analysis of their current situation. Errors at the knowledge based level can also be caused by encysting; operators continue to focus on minute details of a much wider problem.

Reason extends Rasmussen's Skill, Knowledge, Rule distinctions in several ways. In particular, he focuses on the ways in which failures affect all three levels of performance. A distinction is drawn between the error mechanisms that operate before and after the detection of an error. The former include the skill based slips and lapses while rule and knowledge based mistakes, typically, occur after a problem has been identified. Reason also focuses on monitoring failures that prevent an operator from applying effective problem solving techniques at both the knowledge and skill based levels of performance. He argues that skill based behaviour consists of a 'preprogrammed' sequence of operations together with attentional checks that monitor progress towards an objective. The failure of these attentional checks can result in a slip or a lapse. This observation provides GEMS with much of its design power; it may not be possible to eliminate human error but it is possible to improve self-monitoring during task performance. It is also possible to help the detection of potential errors through 'environmental cueing' and the development of appropriate feedback.

The distinctions between these different error mechanisms have helped to guide the investigation of safety-critical incidents. For instance, slips, lapses and mistakes are all included within EUROCONTROL's harmonisation of European Incident Definitions Initiative (HEIDI) for Air Traffic Managment [717]. This project has developed a common vocabulary that can be used to describe the causes of incidents, including human error, across the many different air traffic service providers in European air space. The concepts introduced in the preceding paragraphs are also being widely used in the official reports that are produced in response to accidents and incidents. Without an understanding of the key concepts behind human error, the following excerpt from a recent ATSB investigation would make little sense:

> "The event which precipitated this accident was the unauthorised action of the Train Examiner in moving the points to set the main line for the yard at Ararat. Unsafe acts can take a variety of forms, including absent-minded slips, memory lapses, mistaken intentions and rule violations. Industrial safety studies have indicated that rule violations are frequent contributors to workplace accidents. In most cases, rule violations take the form of well-intended shortcuts which are motivated by a desire to get the job done in a manner that is perceived to be more efficient than that laid down in the rulebook. The action of the Train Examiner in moving the points appears to have been a rule violation, that is, a conscious act which was contrary to procedures. The investigation team was unable to interview the Train Examiner. Nevertheless, the available information suggests that his action was not motivated by any malicious intention. Rather his action appears to have arisen from a desire to assist, combined with a lack of knowledge and experience."
> ([47], page 36)

People continually make mistakes, commit slips or suffer from lapses of attention. Very few errors and violations will ever result in an incident or accident. This apparent paradox is explained by the monitoring activities that were mentioned in previous paragraphs. We regulate our behaviour to reduce the likelihood of an adverse outcome. Occasionally, however, the internal checks and balances will fail. Inattention and fatigue may prevent us from intervening to mitigate the consequences of previous actions. Under such circumstances, we must rely upon the support of automated systems and of other co-workers.

3.6 Team Factors

Previous paragraphs have focussed on individual human error. Little attention has been paid to the problems of coordinating interaction with other members of a working group or team. In contast, Viller [847] provides a summary of social and group performance failures:

- failures due to distraction. These occur when an individual interrupts one of their colleague's tasks. This may be done intentionally where an operator deliberately wants to attract a co-worker's attention. Distractions can also be an unintentional side-effect of one worker's actions on their colleagues.

- failures due to performance effects. Individuals may consistently perform below expectations if they are worried about their actions being monitored or observed by their colleagues. The performance of operators can also be affected if they wish to impress or 'show off' to their colleagues.

- failures that are due to inappropriate human resources in the group. This can occur if there is group members are not competent to perform necessary tasks. It can also occur if there are too many group members who are competent in a small subset of all tasks. In such cases, there can be competition to focus on a few objectives at the cost of other necessary activities.

- socio-motivational failures. There may be 'free-riders' who hide their poor performance by relying on their colleagues in the group. Some operators will carry their colleagues even though they are reducing the effectiveness of the team as a whole. Individuals often feel that it would be disloyal to report the under-performance of their co-workers.

- group coordination failures. The overheads of coordinating group actions can impair the effectiveness of the group as a whole. The division of labour can create bottlenecks where some individuals are forced to wait for considerable periods until their colleagues have completed related activities. In other situation, necessary tasks may be duplicated or omitted because group members failed to understand what the rest of the group expects of them.

- status related failures. Problems can arise if low status group members are discounted or ignored. Conversely, a group may grant undue attention to high ranking individuals. This can be a particular problem where leaders make judgements based on inadequate information or expertise. Incorrect judgements from a high status member can command influence because others respect status rather than the value of the judgement itself.

- group planning and management failures. Groups may create unnecessary sub-tasks. They can also allocate necessary tasks to inappropriate individuals. In either case, the underperformance of key individuals can place additional strain on the group as a whole.

- failures due to inappropriate leadership style. There are two different leadership styles. One focusses on the socio-motivational aspects of leadership while the other focusses more narrowly on 'getting the job done'. An inappropriate balance of either of these styles may jeopardise group success.

- failures due to inappropriate leadership skills. The appointed leader may not have the necessary skills that contribute to both of the roles mentioned above. A lack of appropriate experience or training can leave leader unprepared for the demands that are placed upon them. They are then likely to make decisions that, in turn, make other failures more likely. For instance, they may assign necessary tasks to individuals who are unsuited to those activities. Converely, they may fail to assign key tasks within the necessary timescale.

- failures due to excessive influence of the leader. A high status leader may stifle contrary opinions in situations where they are, themselves, in the wrong. There can be a temporal dimension to this problem if they persist to advocate a policy in the face of adverse evidence. Alternatively, leaders may fail to revise a decision that was initially correct but that was undermined by subsequent changes in their system or the environment.

- failures due to conformity arising from inappropriate informational influence. This occurs when the judgement of one member is based on false evidence or is misunderstood by another member of the group. Some adverse events occur when other failures exacerbate this type of event. For example, an initial failure to understand process data might be compounded by an inappropriate leadership style if others are discouraged from questioning the preliminary interpretation.

- failures due to group polarisation and groupthink. A group may be persuaded by dillusions of its own invulnerability. It may mutually rationalise actions or observations that support the current concensus, it may ignore or discount inconsistent evidence and arguments.

The following incident illustration some of the problems that complicate group work in safety-critical systems. Heathrow air traffic control were using Runway 27 Right (27R) for take off and Runway 27 Left (27L) for landing. There was one Departures officer coordinating traffic leaving from 27R and another Arrivals officer working with aircraft arriving on 27L. The Departures officer was undergoing training with a Mentor. When one aircraft (SAB603) initiated a missed approach. The Departures officer informed the Arrivals officer of a potential conflict with AFR813. , Departures did not inform the Arrivals officer of another aircraft BAW818 that was also taking off at that time:

> "The incident occurred when the weather at LHR (London Heathrow) deteriorated to conditions below that required by SAB (Sabena) 603 on approach. In consequence, the commander initiated a standard missed approach. Air Arrivals saw the aircraft climbing, acknowledged the missed approach to the crew and activated the missed approach alarm. He also informed his colleague, Air Departures, of the manoeuvre and received the information that AFR (Air France) 813 was airborne on a 'Midhurst' SID (Standard Instrument Departure) and that AFR813 would be turned onto a westerly heading. However, he neither saw nor was informed that another aircraft, BAW (British Airways) 818, was also just taking off on a 'Brookmans Park' SID. Based on the information that he had received, Air Arrivals turned SAB603 to the right to achieve maximum separation with AFR813 and also to minimise any disruption to the latter aircraft's flightpath. This resulted in SAB603 and BAW818 coming into close proximity to each other. Air Departures failed to inform Air Arrivals of all the aircraft on departure at the time of the missed approach ecause she did not consider BAW818 as a confliction. This omission was apparently endorsed by the Mentor since he failed to amplify the information passed. Although Air Departures was sitting in the controller's position, the Mentor retained overall responsibility for the duty." [15]

This incident illustrates the dual nature of group interaction in many incidents. On the one hand, the Arrivals and Departures officers created the conditions that led to the incident by failing to ensure that they were both aware of the potential conflicts. On the other hand, effective intervention by the Mentor helped to ensure that an incident did not develop into an accident. The number of failures that are detected and resolved through effective teamwork will far out-strip the number of reported incidents of team-based failure.

It can be difficult for investigators to identify the causes of team-based failures [728]. Many individuals are reluctant to discuss the details actions of their colleagues in the aftermath of an adverse event. Even if it is possible to reconstruct the events leading to a mishap, it can be difficult to understand the reasons why a team acts in a particular way. It is often necessary to understand the complex relationships that exist between the different members of the group in order to explain their interactions. It is possible to ignore some of these problems by viewing team-based incidents as a straightforward extension of single-person failures. For example, Figure 3.5 extends Figure 3.4 to capture the ways in which an individual's cognitive, perceptual and physiological processes might interact with those of their colleagues. The state of the environment is affected by the actions of several operators. These actions can potentially occur at any time during their colleague's activites. Such interventions can hinder, and also support, an individual's situation assessment, planning and action execution. This diagram also illustrates the way in which operators perceive projections of the total state of the system. User 1's view is unlikely to be the same as User 2's and so on. It also reinforces the idea that any group or team 'situation awareness' is likely to be highly distributed. It is not simply based on what each user can observe of their colleague's interventions through their view on some shared state, it is also based on their anticipations and predictions of what their colleagues plan to do. Figure 3.5 is, however, a gross simplification. Group behaviour cannot simply be viewed as the 'sum of its parts'. For example, Kogan and Wallach [452] showed that groups may be more tolerant of risks than the individuals who contribute to a decision. This 'risky shift' has since been question by investigations into teams that seem to be more cautious than their individual members. Myers resolves this apparent paradox by arguing that initial dispositions help to determine subsequent behaviour [556]. If individuals initially favour a low risk solution then the group is liable to urge even more cautious approaches. If individuals initially accept higher risk positions then the group is liable to adopt even higher risk decisions.

Figure 3.5 cannot easily be used to characterise incidents in which teams make inefficient use of the personnel that are available to them. The Transportation Safety Board of Canada provide two

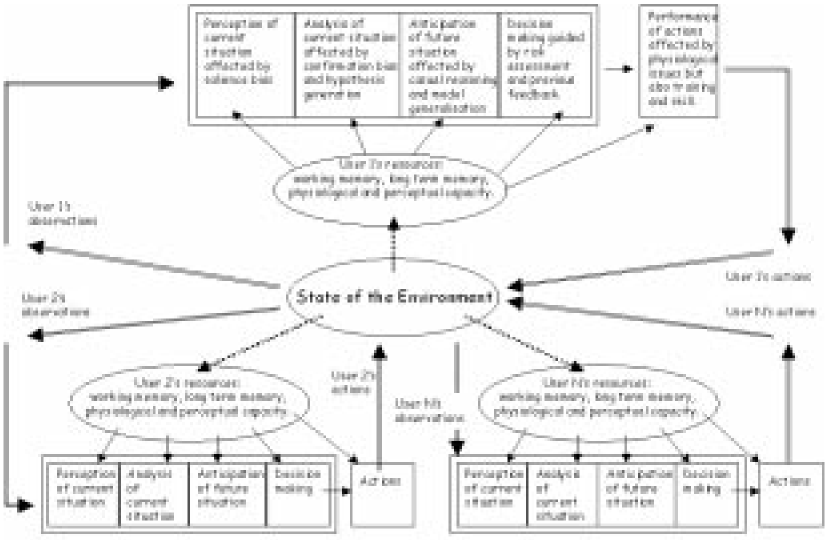
Perception of current situation affected by salience bias
Analysis of current situation affected by confirmation bias and hypothesis generation
Anticipation of future situation affected by causal reasoning and model generalisation
Decision making guided by risk assessment and previous feedback
Performance of actions affected by physiological issues but also training and skill
User 1's resources: working memory, long term memory, physiological and perceptual capacity
User 1's observations
User 2's observations
State of the Environment
User 1's actions
User N's actions
User 2's resources: working memory, long term memory, physiological and perceptual capacity
User 2's actions
User N's observations
User N's resources: working memory, long term memory, physiological and perceptual capacity
Perception of current situation
Analysis of current situation
Anticipation of future situation
Decision making
Actions
Perception of current situation
Analysis of current situation
Anticipation of future situation
Decision making
Actions

Figure 3.5: Cognitive Influences on Group Decision Making and Control

examples involving communication failure between Pilots, Captains and Officers of the Watch:

> "On 08 May 1991, while downbound in the St. Lawrence River with a cargo of oil, the Canadian tanker 'IRVING NORDIC' struck bottom to the north of the ship channel, downstream of the Grondines wharf. The Transportation Safety Board determined that the 'IRVING NORDIC' struck bottom because the vessel left the navigation channel as a result of a premature alteration of course. The alteration of course was ordered by the pilot who believed that the 'IRVING NORDIC' was farther downstream than the vessel really was. The helmsman did not advise the pilot that he was experiencing difficulty in holding the vessel on course. The pilot did not question the helmsman about the position of the wheel relative to the rudder angle indicator. The Officer of the Watch's method of monitoring the vessel's progress was not sufficiently precise to prevent the occurrence. The Board stated that a general lack of interaction and coordination between bridge personnel and the pilot contributed to the accident. (M91L3012)
>
> On 01 July 1991, the loaded Great Lakes bulk carrier 'HALIFAX' grounded in the same area, also due to a premature alteration of course. The Board found that the vessel's position was not double-checked with all available landmarks and navigation aids. The Officer of the Watch was not monitoring the pilot's actions and did not recognize that the change of course was premature. The Officer of the Watch appeared to have placed total confidence in the pilot's navigation ability. When the pilot passed his position report, the Officer of the Watch logged the time, but he did not plot the position on the chart. Had the Officer of the Watch been using a recognized, precise method of monitoring the vessel's progress, he might have been able to recognize the pilot's error and question the change-of-course order before it resulted in the grounding. The Board stated that there was no effective exchange of navigational and operational information (including passage planning) between the officers of the ship and the pilot. (M91L3015)" [619]

Helmreich and Schaffer avoid many of the criticisms that can be made when individual models of cognition, perception and physiology are used to explain the dynamics of group interaction [344]. They provide an alternative view of group interaction in their model of operating room performance. Figure 3.6 is based on this approach. This model has the benefit that is captures many of the sources of failure in the Viller taxonomy [847]. Individual and organisational outcomes are clearly distinguished from those of the team as a whole. The organisational 'culture' and 'norms' are explicitly denoted as contributory factors to group performance. However, it does suffer from some

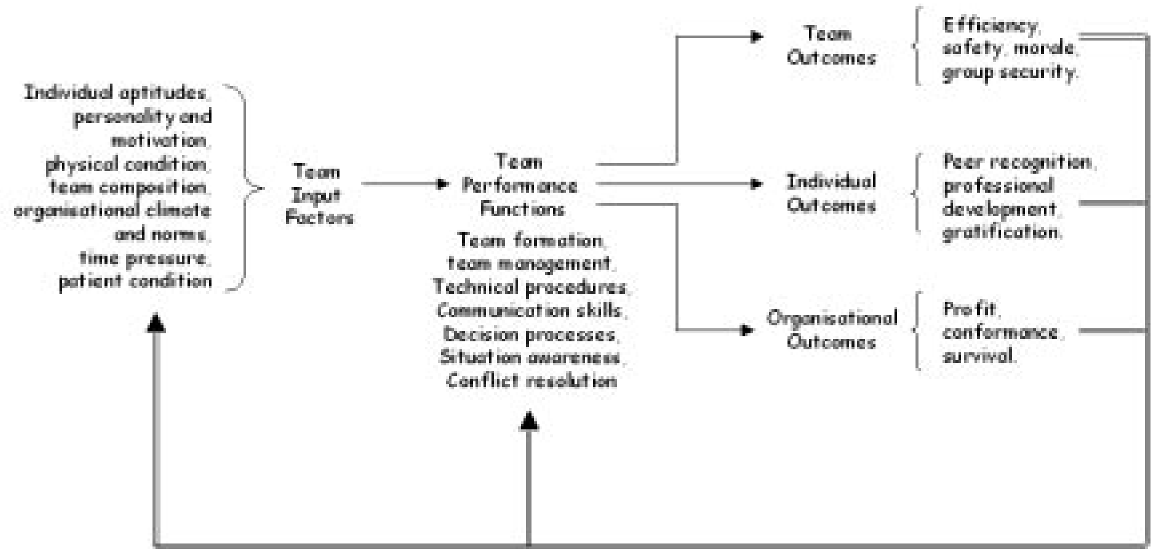

Figure 3.6: Influences on Group Performance

important limitations as a tool for understanding team-based failures. Neither Figure 3.5 nor 3.6 consider the more detailed problems of group-based communication that contribute to most incidents and accidents [64]. This is important because communication failures not only contribute to the causes of an incident but also impair an organisation's ability to respond to the aftermath of an incident. .

3.6.1 Common Ground and Group Communication

Grice [296] has developed a number of guidelines to support communication with groups of co-workers: be as informative as is required but not more so; say what is true, not that for which you lack sufficient evidence; be relevant; be easy to understand, not obscure, ambiguous, verbose, disconnected. A number of authors have identified practical problems in achieving these maxims within many application domains [525]. In particular, it can be difficult to satisfy Grice's maxims when teams must operate under time pressures or under real uncertainty about an individual's understanding of their co-workers' beliefs [168]. In order to understand why it can be difficult to satisfy Grice's guidelines, it is important to undertsand the concept of common ground within group-based communication. A transcript from a cockpit voice recorder can be used to illustrate this point. The account begins imediately before the crew shut-down their one healthy engine:

> "From the Cockpit Voice Recorder it was apparent that the first indication of any problem with the aircraft was as it approached its cleared flight level when, for a brief period, sounds of 'vibration' or 'rattling' could be heard on the flight deck. There was an exclamation and the first officer commented that they had 'GOT A FIRE'. The autopilot disconnect audio warning was then heard, and the first officer stated 'ITS A FIRE COMING THROUGH'. The commander then asked 'WHICH ONE IS IT?', to which the first officer replied, 'ITS THE LE..ITS THE RIGHT ONE'. The commander then said 'OKAY, THROTTLE IT BACK.'
>
> London Air Traffic Control was then called by the first officer, advising them of an emergency, after which the commander asked for the engine to be shut down. The first officer began to read the checklist for 'Engine Failure and Shutdown' but was interrupted by Air Traffic Control calls and the commander's own calls to the operating company during which the decision was made to divert to East Midlands. Approximately 2 minutes after the initial 'vibration' the final command was given to shut down the engine. The

> first officer then recommenced the checklist and 2 minutes 7 seconds after the initial engine problem he moved the start lever of the No 2 engine to ‘OFF’. He then started the APU (Auxilliary Power Unit). Throughout this period no fire audio warning was heard.” [8]

There are several hypotheses about the causes of this error. The events between the crew’s initial conversation and the First Officer’s action might have interfered with the First Officer’s recollection of what had been decided. Alternatively, the First Officer’s comments show some indecision between the Left (No 1) engine and the Right (No 2) engine. This indecision was not reflected in the Commander’s instruction to simply ‘Shut it down’. Clark and Brennan [167] provide means of interpreting such failures. They argue that people are continually trying to ground their conversations. Grounding is the process of seeking and providing evidence of understanding in conversation. This grounding process did not occur in the previous transcript because the Commander believed that the First officer was clear about the source of the problem. The First Officer’s decision to shut down the No. 2 Right engine (and the investigator’s subsequent criticism of the crew’s lack of review prior to this decision) also reflects the way in which the First Officer also assumed that the Captain was sure that the problem lay in the No. 2 engine, in spite of their initial hesitation.

It is important to understand why team members fail to perform the cross-checking that is necessary to ensure they accurately understand the meanings of their colleagues’ utterances. One explanation for this is that establishing common ground will carry a number of potential costs. Table 3.3 lists some of overheads involved in refining our understanding of a converstion. This helps to exlpain why the costs of repairing a potential mistake can be perceived to be more costly than executing an action based on partial knowledge [863]. In other situations, very similar events can lead to entirely different team behaviours. For example, individuals may initiate ask further questions to clarify their understanding of their colleagues’ beliefs and intentions if that indidividual has received appropriate training or if circumstances allow more time for review. In such a situation, the costs of repair may be perceived to be less than the costs of delayed intervention.

Cost	Description
Formulation	formulate and reformulate utterances
Production	producing the utterance
Reception	receiving a message
Understanding	understanding a message
Start-up	starting a new discourse
Delay	planning and revising before execution
Asynchrony	timing of discourse exchanges
Speaker change	changing speakers
Display	presenting an object of the discourse
Fault	producing a mistake
Repair	repairing a mistake

Table 3.3: The Costs of Establishing Common Ground

The likelihood of a fault occuring in the common understanding between operators is heavily influenced by their medium of communication [167]. For example, the time take to repair a mistake will be far greater if the operators are not physically copresent. this may be even greater if temporal distance is also introduced. For example, a common problem in maintenance procedures is to understand the information left about the progress made by previous engineers on previous shifts who may now not be on site:

> “Conscious of the total amount of work which Line Maintenance had to do that night the Line Engineer readily accepted the offer and in the absence of any stage paperwork only gave a verbal handover to the Base Maintenance Controller. Thus he could dispose of the Borescope Inspections and get on with the other Line Engineering work he had

> with minimum delay. He felt that such a brief was adequate as the Base Maintenance Controller was a senior and well respected member of the staff, with the reputation of being highly competent, conscientious and possessing a considerable depth of knowledge of the aircraft types operated by the Company. It was clear from their statements that both the Line Engineer and the Base Maintenance Controller were satisfied, after their verbal exchange, that the existing state of the aircraft and the total requirement of the task were well understood by both.
>
> It is clear, however, from a number of facts revealed during the investigation that the Controller did not fully appreciate what had been, or remained to be, done. He was unaware of the loosened plug, he did not renew the HP rotor drive cover O-rings and he did not complete idle power engine ground runs. " [12]

We have argued that the establishment of common ground is a key objective for team based interaction. We have also argued that many incidents occur because operators fail to ensure that their understanding of their colleagues' beliefs and intentions does reflect those beliefs and intentions. However, it is important to recognise that this only provides a partial accout of team-based failures in incidents and accidents. The previous theoretical work in this area has ignored the ways in which the imperatives of communication change under "adverse" circumstances. For instance, an initial failure to establish common ground may then lead to a situation in which direct orders must be issued and followed without question (or understanding). This is illustrated by the following Air Traffic Control incident involving a Terminal Radar Approach Control (TRACON) team:

> TRACON Supervisor: "Get 487 outta here, send him around"
> Trainee: "I cant - he's changed [his radio] over to the tower"
> [Supervisor reaches between his radar and flight data systems and presses a button that connects him directly with the Local Controller in the tower]
> Local Controller: "Pull United 487 outta here, immediate go around, maintain altitude"
> Local Tower Controller: "United 487 immediate go-around; maintain altitude; maintain runway heading: stay with me".[91]

In the supervisor's view, action was needed immediately without any opportunity to establish the necessary context for the Tower controller to understand the reasons for the order. The Tower controller was prepared to act without stopping to ask about the reason for the message that he had received [91]. On the one hand, such incidents illustrate how key personnel may be trained to act without hesitation if circumstances demand. However, the dangers associated with such actions also illustrate the importance of avoiding these circumstances in the first place.

3.6.2 Situation Awareness and Crew Resource Management

The previous incident shows how communication failures can force individuals to issue 'high-risk' instructions. The trainee failed to directly inform the 487 or the Local Controller of the potential threat before the supervisor intervened. The TRACON supervisor was then forced to issue a 'high-risk' command because they relied upon the Local controller to act without question. However, the key point to understanding this incident is to question why the trainee failed to communicate the potential threat to his colleagues. Many analysis and investigators would asign this to a loss of *situation awareness*. There are numerous definitions of this term [726, 661, 871]. This research work mirrors the numerous phrases that are used to describe the problem in incident report systems: 'falling behind the plane'; 'losing the big picture'; 'spotting the wood for the trees'; 'losing the bubble'. Endsley and Smolensky argue that "situation awareness is the perception of elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future" [225]. They go on to define three levels that contribute to good situation awareness. Level 1 situation awareness consists of the perception of elements in the environment. Level 2 situation awareness focusses on the comprehension of the current situation. Level 3 situation awareness consists of the projection of future states. These distinctions have a great deal in common with the perceptual and cognitive processes illustrated in Figures 3.4 and 3.5.

In contrast, Endsley and Smolensky's distinctions have been used to identify possible causal factors behind incidents reported to the FAA/NASA's Aviation Safety Reporting System [432].

This study focussed on 33 incidents of poor situation awareness in Air Traffic Control. 69% involved failures at level one, 19% involved failures at level two, 12% involved failures at level three. Such ratios should not be surprising given that a failure at level one is hardly likely to support adequate performance at level two or three. Of the level one failures, the loss of situation awareness was most often due to a failure to monitor or observe data (51.5%). Most of these incidents were caused by distractions (53%), high workload (17.6%) and poor vigilance (11.8%). Later sections of this book will describe the problems in replicating these subjective classifications. For now, however, it is sufficient to observe the paradox that often arises in detailed studies of situation awareness. Problems in our perception of our environment, typically, stem from unnecessary signals or interruptions from that environment. In other words, incidents are often caused by disruptions that are created when information is presented to us that might, in other circumstances, have been essential to our control tasks.

At the heart of situation awareness problems is the difficulty of monitoring mutiple, simultaneous processes. This problem has particular relevance for team based interaction because, as noted in the previous paragraph, inefficient group communications jeopardise successful anticipation of future states. This is illustrated in the following report:

> "The CVR transcript reveals that the flight engineer was overloaded and distracted from his attempts to accomplish the Fire & Smoke and Cabin Cargo Smoke Light Illuminated emergency checklists (in addition to his normal descent and before-landing checklist duties) by his repeatedly asking for the three-letter identifier for Stewart so that he could obtain runway data for that airport.
>
> The captain did not call for any checklists to address the smoke emergency, which was contrary to FedEx procedures. Nor did he explicitly assign specific duties to each of the crewmembers. The captain also did not recognize the flight engineers failure to accomplish required checklist items, provide the flight engineer with effective assistance, or intervene to adjust or prioritize his workload. In fact, the captain repeatedly interrupted the flight engineer during his attempts to complete the Fire & Smoke checklist, thereby distracting him further from those duties.
>
> The Safety Board concludes that the captain did not adequately manage his crew resources when he failed to call for checklists or to monitor and facilitate the accomplishment of required checklist items. Therefore, the Safety Board believes that the FAA should require the principal operations inspector for FedEx to review the crews actions on the accident flight and evaluate those actions in the context of FedEx emergency procedures and training (including procedures and training in crew resource management) to determine whether any changes are required in FedEx procedures and training." [591]."

The previous report is interesting for a number of reasons. Firstly, it shows how team based interaction is often critical in the aftermath of an incident. The crew were one of the key defence mechanisms for the system once the initial fire had taken hold. Secondly, as noted above, it illustrates how inefficient leadership and task allocation can jeopardise the coordination that is necessary in extreme circumstances. Finally, the Safety Board illustrate how "procedures and training in crew resource management" are perceived to support crew coordination during adverse circumstances.

Crew Resource Management (CRM) techniques have been developed to improve group coordination during incidents and accidents [733]. A number of recommended practices have been introduced into the aviation and maritime industries to encourage mutual situation awareness, team-based decision making and workload management. Initially, these practices focussed on an individual's interaction with their colleagues [343]. Training materials focussed on the use of protocols and procedures that reduced ambiguity in crew communications. They, therefore, owed more to the Gricean maxims than Clark's emphasis on an iterative search for common ground. More recently, CRM training has focussed on team building and the effective sharing of tasks during high-workload situations [91]. This was reflected by a change in the use of terms such as "cockpit resource management" to the more general "crew resource management". This has reached the point were current

CRM techniqus also consider the role of ground staff and of cabin crew during incidents and accidents. CRM training is now a pre-requisite for public transport operators to be granted their UK Aircraft Operators Certificate. UK Aeronautical Information Circular 143/1993 states that all crew must have been trained in the importance of Standard Operating Procedures, the Flight Deck Social Structure and a detailed examination of the manner in which CRM can be employed in order to make a positive contribution to flight deck operations. Joint Airworthiness Requirements (JAR OPS) sub-part N, 1.945(a)(10) and 1.955(b)(6) and 1.965(e) extended similar requirements to all signatory states during 1998. Similar initiatives have been proposed for maritime regulations:

> "On June 25, 1993, as a result of its investigation of the grounding of the United Kingdom passenger vessel RMS Queen Elizabeth 2 (near Cuttyhunk Island, Vineyard Sound, Massachusetts, on August 7, 1992, the Safety Board issued Safety Recommendations M-93-18 and -19 to the Coast Guard. The Safety Board requested that the Coast Guard: Propose to the International Maritime Organisation (IMO) that standards and curricula be developed for bridge resource management training for the masters, deck officers, and pilots of ocean-going ships. (M-93-18) Propose to the IMO that the masters, deck officers, and pilots of ocean-going ships be required to successfully complete initial and recurrent training in bridge resource management. (M-93-19)
>
> As a result of its investigation of this accident (grounding of Panamanian Passenger Ship, the Royal Majesty), the NTSB reiterates the following recommendations:
>
> To the U.S. Coast Guard: Propose to the IMO that standards and curricula be developed for bridge resource management training for the masters, deck officers, and pilots of ocean-going ships. (M-93-18) Propose to the IMO that the masters, deck officers, and pilots of ocean-going ships be required to successfully complete initial and recurrent training in bridge resource management. (M-93-19)" [594]

The IMO's Sub-Committee on Flag State Implementation and its working group on casualty statistics and investigation continue to show an active interest in following the legislative and regulatory lead established by the JAR OPS provisions, mentioned above.

It is possible to identify two different approaches to the use of modern CRM training. Firstly, CRM training is used to support crew coordination under those rare emergency situations that impose the greatest workload [91]. High-fidelity simulators are used to help crews test team-performance in a direct manner. This approach is widely associated with Foushee and Helmreich [279]. In contrast, the second approach rejects this focus on the simulation of extreme situations. Seamster and others [733] have argued that crew coordination practices are ingrained more deeply if they are treated as a key component of many routine tasks [734]. It is important to note that these two approaches need not be contradictory. Simulator training may also be used to back-up more routine applications of CRM training. The difference lies in the emphasis that Seamster and others have placed upon the use of CRM techniques in nominal operating conditions. However, incident reporting schemes introduce a filter or bias. Submissions are more likely to report extreme forms of good CRM than more everyday instances of appropriate behaviour. For instance, the following excerpt shows how extreme circumstances force a crew to simultaneuosly address a number of failures that could not easily have been predicted or anticipated before the incident itself.

> "The Captain's autopilot dropped off with several warning flags on his flight instruments. He transferred control of the aircraft to me. During descent, various warning lights illuminated, which were reset several times. We ended up with one pitch trim working. The Captain was surrounded by inop flags on his instrument panel, so was unsure of which instruments were still operating. Random electrical warnings erroneously indicated that the aircraft was simultaneously on the ground and in the air. The Captain and I had donned oxygen masks as soon as we detected smoke. The Captain had a partial com. failure with his oxygen mask, then with his headset/boom mike. Cabin pressurization was climbing.
>
> Cabin pressurization control was switched to standby mode. The Second Officer found a second fire extinguisher and discharged it into the continuing red glow in the

> circuit breaker panel. During the approach, we encountered... failure of both direct lift control auto spoilers. At touchdown, spoilers were manually extended. I selected reverse thrust, but no thrust reversers worked. On taxi in, all three engines were in flight idle. At the gate... the aircraft was still pressurized. Flight Attendants could not open the door.
>
> The Second Officer tried to shut down all packs and engine bleeds, but could not. The Captain attempted to shut down the engines with fuel and ignition switches, but engines kept running. Engine fire [fuel shutoff] handles were pulled, and engines shut down. The door was opened from the outside, and the passengers exited.
>
> [Comment from ASRS editors] The final diagnosis from maintenance personnel: an improperly installed wiring clamp had worn through the insulation and shorted out. Kudos to the flight crew for great crew coordination and superb handling of this aircraft emergency." [57]

The previous example is clearly an unusual incident. The nature and extent of the systems failure forced the crew to take relatively extreme measures, such as discharging a fire extinguisher into a circuit breaker panel. This incident is also atypical in that it focusses quite narrowly on the coordination between members of the flight crew. It ignores wider forms of cooperation that typify many safety critical systems. The working group of a pilot and co-pilot clearly extend well beyond the flight deck to include cabin crew, air traffic control etc. The following report from the Aviation Safety Reporting System illustrates this more general aspect of appropriate CRM behaviour:

> "Some reporters continued with an operation even when something didn't look right, or was blatantly wrong. Flight crews also admitted to failing to request a tug to get into, or out of, a tight parking place. The latter two problems may have been responses to schedule pressure or to demand for on-time performance, also mentioned by many flight crew members as an underlying cause of incidents. These and other sources of distraction also caused a marked reduction of cockpit coordination and CRM skills. A plane's rear airstairs received damage when the crew became distracted by multiple demands, and failed to act as a team:
>
> "[This incident was caused by] distractions in the cockpit, plus a desire to operate on schedule. There were several conversations going on from inside and outside the aircraft.
>
> Raising the airstairs is a checklist item... backup is another checklist item which requires the Second Officer to check a warning light. No one noticed the light. The pushback crew consisted of 2 wing observers plus the individual in the tug...all failed to observe the rear stairs." [159]

Previous paragraphs have argued that CRM techniques can be used to address some of the team-based failures that are identified by incident reporting systems. Later sections will go on to show how incident reporting systems can be used, arguably for the first time, to question the success of such techniques. For now, however, it is sufficient to observe that good CRM is no guarantee of good team interaction. Training alone cannot easily counteract some of the social and leadership issues that were identified in Viller's list of the causes of team failure [847]. For example, a recent NASA Ames study reinforced many informal observations from incident reports when it concluded that Captains tend to be pro-active in high-risk situations; often preventing these situations from developing through pre-emptive actions. First officers were sensitive to the social dynamic of challenging the captain. They were most likely to intervene in situations involving *external* errors when risk levels were high [662].

3.7 Summary

This chapter has summarised the factors that contribute to incidents in safety-critical applications. Many stem from regulatory failures. For example, regulators have ignored, postponed and only partially implemented the recommendations from previous incidents only to find that they recurr a short time after the initial occurence. With limited resources, it is difficult for such national and regional organisations to effectively monitor increasing complex, heterogenous production processes. This has created a situation in which regulators are dependent upon information from line-management. This information increasingly comes through participation in national and international incident reporting schemes.

Incidents also occur because managers fail to recognise or satisfy their regulatory obligations. They can occur if management fails to perform the usual leadership functions that are expected in safety-critical industries. For instance, managers may fail to support an adequate safety-culture. It is important not to underestimate the practical difficulties of avoiding such failures. It is notoriously difficult to identify quantitative measures for the success or failure of such management objectives. The visible attributes that are associated with a good 'safety culture', such as the maintenance of an incident reporting scheme, often reflect a desire to conform with regulatory requirements rather than a pro-active attitude to the prevention and mitigation of accidents [673]. Even where safety-culture is supported, it can be difficult for managers to ensure that best practice propagates throughout large, complex and dynamic organisations.

Management failures helps to establish the latent conditions for future incidents. For example, inadequate maintenance schedules contribute to more catalytic hardware failures. Decisions to sacrifice redundant protection devices leave systems vulnerable to transient faults. These examples illustrate how concern is incresingly focussing on these more organisational aspects of hardware failure: in acquistion; in testing and validation and in maintenance scheduling. Many of the more technical aspects of hardware reliability are now well supported through the provision of appropriate tools ranging from application specific CAD/CAM environments through reliability methods, such as Failure Modes, Effects and Critical Analysis, to more abstract mathematical techniques, such as Markov Modelling and Monte Carlo simulation. It is, therefore, not surprising that incident reporting systems have long been used to support the acquisition and validation of hardware reliability data, for instance through the Failure Reporting, Analysis and Corrective Actions (FRACAS) schemes advocated by the US Department of Defense.

Software failure pose an increasinly important challenge for the management of safety-critical systems. The probabilistic techniques that can be used to assess and predict hardware failure rates cannot easily be used to analyse the reliability of software systems. The lack of what we have termed 'forensic software engineering' techniques also leave us vulnerable to repeated failures. In particular, recent investigations of accident and incident reports has revealed a number of technical and pragmatic concerns that limit the recommendations of many investigations. The current focus on process based standards for software development creates further challenges. Incidents of software failure raise doubts not simply about the quality of certain modules and procedures or about the ability of individual programmers. Such failures bring into question all of the code that has been produced using that particular development process.

Human-computer interfaces represent one of the key areas in which software contributes to the causes, or exacerbates the consequences, of safety-critical incidents. Such interaction problems stem from a complex blend of design failures, of incompatabilities between the tool and its context of use and of human 'failure' [126, 125]. Several taxonomies have been developed to help analysts categorise the different forms of human error and violation that jeopardise system safety. These taxonomies provide convenient labels for talking about the human contribution to incidents. Unfortunately, many incident reporting schemes simply record frequency data for each of these categories. It is important to go beyond terms, such as slips and lapse, to understand the perceptual, cognitive and physiological per-cursors to errors and mistakes. It is also important to understand the ways in which individual characteristics and social pressures contribute to the necessary conditions for failure. Conversely, however, it is important to recognise that operators resolve many situations that might otherwise have resulted in incident or accident reports. There is a danger that the analysis of human error will mask instances in which human intervention preserves the safety of application processes.

Many incidents are caused not simply by individual instances of human failure but by the problems of group decision making. Some of these problems stem from organisational problems. It can be difficult to identify an efficiently allocation of shared tasks to the members of a team. It can be difficult to identofy individuals with the necessary leadership skills and so on. Other problems relate more narrowly to issues of group communication. Under stressful situiations it can be difficult to ensure that the members of a group know about not just current actions of their colleagues but also their future goals and intentions. Without some shared understanding of this information then the situation awareness of each member of the group is liable to be compromised. As with the other causes of safety-critical incidents, group failures also raise important problems for the establishment and maintenance of incident reporting systems. It can be very difficult to reconstruct a coherent account of many incidents given that the different individuals in a group are liable to share different understandings of the events leading to failure.

The previous paragraphs have, to some extent, introduced false distinctions betweem regulatory failure and managerial weakness, between hardware failure and software problems, beween individual human failures and team-based failures. This has been a considerable weakness both of existing incident reporting schemes and of academic research in this area. Too many models and techniques focus on specific causal factors. For instance, human error models often concentrate on the phenotypes of inidividual performance without providing any guidance or analytical power for team-based failures. Conversely, techniques for requirements engineering that can be applied to represent and reason about the causes of software bugs often cannot be applied to analyse regulatory failure. The intention of this book is to break down some of these distinctions and and the same time to illustrate both the strengths and weaknesses of many of the techniques that have traditionally support incident analysis. The primary means of achieving this is to continually refer to the complex, pathological events that contribute to real incidents. The strengths of existing models are demonstrated by the analytical insights that they yield into particular instances of failure. Their weaknesses are demonstrated by the ways in which they can obscure or ignore other contributory causes. Before we can extend this investigation of analytical techniques, it is important first to look at the ways in which we can elicit information about safety-critical incidents.

Chapter 4

The Anatomy of Incident Reporting

The following incident report was recently published by the Australian Transportation Safety Board (ATSB). It describes an incident that was initially notified by a member of the public and which was subsequently investigated by ATSB staff:

> "A member of the public reported seeing a single engine aircraft manouevre suddenly to avoid another aircraft, on an intersecting track, while the aircraft were over Brisbane.
>
> An investigation reviewed radar data and air traffic control automatic voice recordings to establish the sequence of events. The investigation found that VH-OXF, a Beech 300, was tracking for a left base to runway 01 at Brisbane Airport at 2,500 ft, while a Cessna 172, VH-IGA, was tracking north over the suburbs at 1,500 ft. The Brisbane departures controller established that the pilot of the Beech could see and was able to avoid the Cessna before reducing the vertical spacing between the aircraft to less than the vertical separation standard of 1,000 ft. The Beech pilot reported seeing and passing over the top of the Cessna and ready for further descent. The controller issued a clearance for a visual approach. The recorded radar data indicated that the Beech began a steady descent from about the intersection of the aircraft tracks.
>
> The controller's options in relation to ensuring separation between the aircraft were either to: maintain the Beech at 2,500 ft until there was more than 3 NM lateral separation with the Cessna; or use visual separation procedures by having a pilot report seeing the other aircraft and then instructing that pilot to avoid the sighted aircraft. To enable the Beech to descend in preparation for landing, the controller used the second option. Examination of the radar data indicated there was no infringement of separation standards.
>
> The recorded radar data indicated that during the period when the Beech was assigned 2,500 ft, the Mode C altitude intermittently indicated 2,300 ft and 2,400 ft. Mode C altitude has a tolerance of plus or minus 200 ft. The pilot was therefore complying with the air traffic control clearance." [50]

This report illustrates some of the tasks that must be performed during any incident investigation. The incident must be reported to the appropriate authorities. The people who initially receive a notification must take any immediate action and pass it on for further investigation. The conclusions and findings of any investigation must be published. Although there were no immediate recommendations from the incident cited above, if there had been then these must be implemented and monitored. The following sections take each of these tasks or roles and considers how they contribute to the successful implementation of an incident reporting system.

4.1 Different Roles

It is important to emphasise that the following paragraphs identify tasks or roles. Any individual or group can perform several of these roles depending on the organisational needs of the reporting system. For example, in a local system the same individual may both receive a report and conduct any subsequent analysis or investigation. In a national or international system, it is more likely that specialist analytical expertise might be called upon to support the local officers who receive an initial notification.

Section 4.2 builds on this analysis and describes a number of different organisational models that can be used to manage these different roles or tasks within specific working environments.

4.1.1 Reporters

This is the person who contributes the initial incident report that triggers any occurrence investigation. The organisation running the scheme may employ them, they may be subcontractors or they may be employed by other organisations that must co-operate with the organisation running the scheme. For example, a member of an aircrew might report an air traffic control incident. Alternatively, members of the public who have witnessed or been involved in an occurrence report some incidents, as was the case in the incident cited above. The following paragraphs identify some of the issues that must be considered when encouraging such contributions to occurrence reporting schemes.

Am I Protected?

Previous chapters have argued that incident reporting systems depend upon the trust of those who contribute to them. If individuals are concerned about punitive actions or about the confidentiality of their submissions then they are unlikely to participate in such a system. One means of preserving this sense of trust is to publish a summary of the rights that protect workers who contribute to a reporting scheme. These rights are partly built on legislative protection, they also rely upon the procedural safeguards that support their participation during the investigation and analysis of an incident.

It is important that the individuals who contribute to an incident reporting system are aware both of their rights and responsibilities when contributing information about adverse occurrences. For instance, in some industries it may be assumed that operators have the right to be excused from further duties in the aftermath of an incident until they are physically or psychologically fit. It is important that such actions should not be interpreted as an admission of guilt or responsibility for an incident. Some systems also offer various forms of counselling to support individuals int he aftermath of an adverse occurrence.

Operators often have the right to a representative of their choice during subsequent interviews or hearings. These representatives can be colleagues, lawyers or a trades union officials. Their presence can have a profound impact both on the individual's participation in a system but also on wider perceptions about the efficacy of incident reporting. There are also practical implications. It can be difficult to schedule investigatory meetings if workers' representatives are unavailable when investigators must compile evidence about an occurrence.

Many national legal systems preserve an individual's right to silence during criminal investigations. Incident reporting systems are not concerned with such criminal acts. However, many systems do offer individuals the opportunity not to 'incriminate' themselves. Operators are not obliged to make written statements. Other systems do not go this far but do ensure that individuals can consult with their chosen representatives before submitting written material.

After the initial information has been gathered about an incident, it is important that workers are aware of their rights during any subsequent analysis. For instance, workers and their representatives may have the right to pose questions to the investigation team. They may also be entitled to review any relevant documents, data recordings or transcripts before appearing in front of any enquiry. Finally, it is also possible for contributors to review the contents of a final report and offer a written response that may be appended to the initial document.

It can, of course, be argued that these various arrangements add greater administrative complexity to incident investigation. Worker representation and participation may also 'tie the hands' of incident investigators. However, such arguments must be balanced against the primary importance of ensuring participation and consensus. Unless individuals feel confident of equitable treatment then they will not contribute. Unless groups of workers are confident in the findings of an investigation then they may oppose the implementation of controversial findings and recommendations.

Should I Report?

A key issue here is that potential contributors must know about the scheme and know how to submit a report. The scale of this task can be illustrated by the distribution list associated with the UK Medical Devices Agency's (MDA) reporting scheme for Adverse Incidents and Disseminating Safety Warnings. This list describes those who must pass on information about this scheme to the people on a far larger list of potential contributors:

> "Please bring this notice to the attention of all who need to know or be aware of it. This will include distribution by:
>
> TRUSTS to:
> Liaison Officers (for onward distribution), All relevant staff including: Risk Managers, Safety Officers, Medical Directors, Clinical Directors, Nurse Executive Directors, Medical, Dental and Nursing staff, Medical Physics/EBME, Operating Theatres, Intensive Care Units, Intensive Therapy Units, Ambulance staff and Paramedics.
> HEALTH AUTHORITIES to:
> Liaison Officers (for onward distribution), Chief Executives of Primary, Care Groups, Registration Inspection Units, General Medical Practitioners, General Dental Practitioners, Opticians, Pharmacists, Practice Nurses, Nursing Homes, Hospices, Private Hospitals.
> SOCIAL SERVICES to:
> Liaison Officers (for onward distribution), Registration Inspection Units, Residential Care Homes, Occupational Therapists, Special Schools." [535]

The scale of this task should not be underestimated. These distributors must ensure that induction courses and periodic retraining reminds staff about the importance of reporting. They must also perform more prosaic duties. For example, they must ensure that staff are providing with access to reporting forms at all times. The logistics involved in disseminating information about incident reporting systems are not the only challenge

It is not enough simply to inform potential respondents about reporting procedures, they must also be able to provide the necessary details that are requested by forms or other elicitation documents. This is a non-trivial task. it can be difficult to draft a form that will elicit sufficient information from all of the many different groups listed above. If respondents do not understand a question then they may fail to provide necessary information. If they misinterpret a question then they may provide erroneous or misleading responses. All of these issues have been compounded by the increasing use of electronic submission forms based on Internet technology. The design of these submission procedures will be discussed in greater detail in Chapter 5.

Will Everyone Participate?

This will be a continuing theme throughout much of this book. Previous chapters have cited the relatively low participation rate in voluntary aviation reporting schemes by general aviation and the military in contrast to commercial aviation. It should also be noted that such comparisons are compounded by the difficulty of estimating what the anticipated reporting rate ought to be. It can be difficult to assess whether each of these groups has a comparative exposure to hazardous occurrences etc. For example, in one local incident reporting system within a UK intensive care unit, approximately 90% of all reports were submitted by nursing staff over a ten year period. 621 reports were submitted by nurses compared with 77 reports by medical staff. However, these figures must be interpreted in terms of the number of staff on the ward. Usually the team consisted of three medical staff, one consultant, and up to eight nurses per shift. The analysis is further complicated by the fact that nursing staff had the most direct contact with the patients who remain the focus of the reporting system and hence may have had proportionately greater opportunity to witness adverse events [119]. Each of these factors must be considered before concluding that there is a systematic under-reporting by medical staff.

What Did I Really See?

There are clear problems in interpreting the evidence provided by an initial report of an incident. For example, the testimony of one eye witness to the Concorde crash was initially interpreted an being consistent with the illumination caused by afterburners rather than a fire involving the fuel tanks. Statements that indicated the true extent of the damage to the aircraft on take-off were dismissed as the exaggerated claims of uninformed observers. The problems of interpreting eye witness statements are not simply related to the difficulty of assessing non-technical accounts of system failures. They can also arise when qualified personnel attempt to provide immediate causal explanations. As mentioned in previous chapters, witnessing an accident can often have the effect of confirming previous concerns about particular operational problems. This confirmation bias can dissuade technical witnesses from considering alternative hypotheses in the immediate aftermath of an incident or accident. A feeling of direct personal responsibility or of physical threat during an accident can lead witnesses to either minimise of maximise the implications of the incidents that they report. Conversely, as mentioned in previous chapters, reports may be contributed by individuals who are more concerned with a perceived grievance than with the overall objectives of addressing safety issues. Reports may also be biased in order to protect themselves, their co-workers or their employers. This final point is illustrated by the findings of an enquiry into a trench collapse that was reported by the US Occupational Safety and Health Administration (OSHA). This refers to witness testimony in the investigation of an incident rather than the initial report of an incident. However, the following quotation does illustrate the potential problems of interpreting bias in eye-witness statements:

> "The judge based his finding that the trench walls had no significant slope on the testimony of 'three disinterested on-the-scene eyewitnesses' (two paramedics and a volunteer fireman), who entered the trench that collapsed. All three reported seeing identical conditions. The judge found the testimony of two corporate officers and two other employees of Zunker regarding the trench dimensions and sloping to be 'unreliable and indeed untruthful,' stating as follows:
>
> The testimony of all these witnesses, each of whom had an interest in the results of these proceedings, was at total odds with the testimony of the [paramedics and fireman] who were disinterested and who truthfully reported their observations at the work site, and in particular at the site of the cave-in. The demeanor of [Zunker's witnesses] as well as their sworn testimony, leaves much to be desired as having any probative value in determining the factual issues in this case.... What element of truth we do attribute to these witnesses comes from Respondent's backhoe operator who indicated that it took him 20 minutes to dig the trench.... [I]t would be virtually impossible to excavate a trench in accordance with the dimensions testified to by [Zunker's president] within a 20-minute period." [646]

The problems that arise immediately after the reporting of an incident are compounded in anonymous systems. This is best illustrated by US guidelines that provide recommended practices for small businesses following the notification of any incident:

> "Gather evidence from many sources during an investigation. Get information from witnesses and reports as well as by observation. Interview witnesses as soon as possible after an accident. Inspect the accident site before any changes occur. Take photographs and make sketches of the accident scene. Record all pertinent data on maps. Get copies of all reports. Documents containing normal operating procedures, flow diagrams,

> maintenance charts, or reports of difficulties or abnormalities are particularly useful. Keep complete and accurate notes in a bound notebook. Record pre-accident conditions, the accident sequence, and post-accident conditions. In addition, document the location of victims, witnesses, machinery, energy sources, and hazardous materials." [649]

These guidelines are generic; they are applicable to a wide range of industries. They also cover what we have termed 'local' reporting systems because they are specifically intended for small businesses. However, these guidelines also illustrate the problems of responding to an anonymous report of an incident. It can be difficult to know where to begin gathering evidence if a report is anonymous. As mentioned in previous chapters, this initial investigation may itself be enough to sacrifice the trust of the contributor and compromise their anonymity. Without the active participation of a known reporter it can be difficult to obtain the additional information that may be necessary to accurately record pre-accident conditions, as they saw them.

Chapter 5 will address these concerns in greater detail. In contrast, the following paragraphs look beyond those individuals who contribute occurrence reports to look at the people who must initially respond to their notifications.

4.1.2 Initial Receivers

The reporter sends their submission to an 'initial receiver'. In most company's incident reports are made to line supervisors unless they are directly implicated in an incident. This has the advantage that supervisors will be familiar with working practices and can take immediate remedial actions to mitigate any adverse consequences. However, these initial receivers need not be directly connected with the reporter's organisation or company. In particular, national systems often rely upon independent reporting agencies. For example, NASA is responsible for administering the Aviation Safety Reporting System (ASRS) on behalf of the FAA. Such organisations protect the notifier's anonymity whilst still enabling investigators to perform subsequent data collection.

The receivers of an incident report are responsible for making an initial criticality assessment. 'Triage' is an important task during the operations of many incident reporting systems. The individual who detects a potential incident must first decide whether or not it is worth reporting. The individual who receives an initial report must then decide whether to pass it on. If they decide that it should be acted on then they must determine who is best placed to act on the report. The group or individual who must act on a report has further more detailed technical judgements to make about the best way in which to investigate and resolve any safety concerns. These decisions depend upon assessments of the importance or priority associated with an incident. Such assessments must be documented and justified in order to support the external inspections that help to ensure consistent responses to similar incidents. The initial receiver also plays an important role in taking any immediate actions that is necessary to safeguard services following an incident. The following paragraphs consider these issues in more detail.

How to Safeguard the System?

The most important task facing the individual who receives an incident report is to coordinate the immediate response to an adverse occurrence. Typically, such actions cannot be delayed until after a full investigation has been instigated or a final report has been delivered. Operators may have to be removed from their working positions. Faulty equipment must be disconnected. Alternative systems or manual back-ups must be set-up. All of this relies upon individuals making a rapid assessment of the context in which an incident occurred. It also relies upon their ability and willingness to instigate immediate corrective actions. Such a response relies upon both a number of factors. The individuals who assume this role must be training to enable them to perform an accurate initial response. They must be familiar with the relevant procedures involved in instigating immediate corrective actions and must feel comfortable with the responsibilities that are associated with such actions. There are clear safety implications if these individuals feel that they lack the appropriate authority or responsibility for taking immediate corrective actions.

It is important to emphasise that operators should not, typically, be withdrawn from their working positions for disciplinary reasons in the aftermath of an occurrence. This would create a strong disincentive to further participation in any such system. In contrast, the purpose behind their removal is to act in the operator's own interest and to preserve the continued safety of their system. In some industries, the knock-on effects of such actions may have relatively minor implications for the operation of the system as a whole. In other industries, the removal of key personnel can impose considerable practical burdens on their colleagues who must continue to operate their systems. These problems are likely to be exacerbated by the fact that many incidents occur during periods of peak workload. As a result, incident reporting systems are often integrated into more general techniques for contingency planning during safety-related failures. The removal of key members of staff can, of course, have further safety implications if their replacements are less well trained, fatigued or nervous about stepping into their roles in the aftermath of an incident. Irrespective of the immediate decision, it must also be determined whether or not an operator should be allowed to return to normal operations or should be relieved for an extended period. This decision has important implications if an investigation determines that inadequate training was a contributory factor to any incident. Clearly such decisions should not be devolved to the person receiving an initial report but must be the shared responsibility of operational and safety managers within the organisation.

It may not be possible for operators to be removed from their duties in confidential or anonymous schemes without raising the suspicion of their colleagues and supervisors. In an open system, however, the removal of staff involved in an incident helps to reduce the likelihood of further failures in the aftermath of an adverse occurrence. It also provide a number of additional benefits. For instance, it can provide an opportunity for those individuals to complete reporting forms while the details of an incident are still 'fresh' in their mind. It also creates an opportunity for the stress management and peer counselling services that are increasingly being introduced in safety-critical industries. These activities are intended to combat the sense of guilt and blame that often arise in the aftermath of an incident. Wu provides a direct impression of the problems that these feelings can cause in the medical domain:

> "In the absence of mechanisms for healing, physicians find dysfunctional ways to protect themselves. They often respond to their own mistakes with anger and projection of blame, and may act defensively or callously and blame or scold the patient or other members of the healthcare team. Distress escalates in the face of a malpractice suit. In the long run some physicians are deeply wounded, lose their nerve, burn out, or seek solace in alcohol or drugs. My observation is that this number includes some of our most reflective and sensitive colleagues, perhaps most susceptible to injury from their own mistakes.
>
> What should we do when a colleague makes a mistake? How would we like others to react to our mistakes? How can we make it feel safe to talk about mistakes? In the case of an individual colleague it is important to encourage a description of what happened, and to begin by accepting this assessment and not minimising the importance of the mistake. Disclosing one's own experience of mistakes can reduce the colleague's sense of isolation. It is helpful to ask about and acknowledge the emotional impact of the mistake and ask how the colleague is coping." [877]

Such counselling helps to maintain valuable human resources, for example, by reducing the likelihood of needing the additional costs of staff replacement. Many organisations provide these services through a peers group who are chosen by the workers themselves and who complete an appropriate training course.

Is the Report Relevant?

As mentioned above, the person or group who initially receives an incident report must determine whether or not the incident falls with in the scope of the system Two different sets of problems are created depending on whether the incident is considered 'appropriate' or not.

If the initial receivers of an incident report believe that it does not fall within the scope of the system then they must reject it. This creates the possibility that important lessons about previous failures will be excluded from the system. In national and international systems, it is also possible that different regional definitions of relevance will lead to inconsistency and bias in the information that is collected. As a result, many organisations publish exhaustive lists of those sorts of incidents that fall within the scope of the system. Some of these lists were considered in the opening chapter of this book when it was argued that it can be extremely difficult to support such exhaustive definitions in complex and dynamic industries where the nature of those failure that are observed will change over time. The problems of determining whether or not an incident falls within the scope of the system are not simply related to technological change. They also relate to the political and organisational environment that support the reporting system. For example, the US Federal Railways Administration published the following exemptions in response to industry objections to the burdens imposed by an occupational injury reporting system:

> Partial relief to certain small railroads generally covered by Part 225. FRA recognises that small operations are concerned with the burdens, both in terms of time and expense, associated with full implementation of the amendments to Part 225 issued in 1996. Based on additional analyses, FRA concludes that it can grant partial relief to certain small operations without compromising the accuracy of its accident reporting data base. These operations are: 1. Railroads that operate or own track on the general railroad system of transportation that have 15 or fewer employees covered by the hours of service law ... and 2. Railroads that operate or own track exclusively off the general system... If your railroad is subject to Part 225 at all and falls in either of the above categories, then you need not adopt and comply with components 3 through 10 of the Internal Control Plan requirements in Section 225.33. See Section 225.33(a)(3)-(10). However, you must fulfill the requirements of components 1 and 2, which require a stated policy dealing with harassment and intimidation. See Section 225.33(a)(1)-(2). To assist railroads in developing this policy, the FRA has provided suggested language, found in Appendix I to this Guide, that may be used. A railroad in either of these two categories is also exempted from the requirements in Section 225.25(a)-(g) to record accountable injuries and illnesses and accountable rail equipment accidents. (See Chapter 2 for definition of accountable events.) You must also, however, maintain a Railroad Employee Injury and/or Illness Record of any reportable condition of one of your employees. (See Chapter 4.) Additionally, a railroad that is generally subject to Part 225 but that operates exclusively off the general system (including off-the-general-system museum and tourist railroads) is not required to report or record an injury or illness of any person that results from a non-train incident, unless the non-train incident involves in-service railroad equipment. See definition of non-train incident in Chapter 2. Railroads that are subject to Part 225 in the first place and that operate exclusively off the general system must, however, continue to comply with Part 225 requirements regarding reporting and recording injuries and illnesses incurred by any person that result from a train accident, train incident, or a small subset of non-train incidents that involve railroad equipment in operation but not moving." [233]

If the person receiving a report can interpret such exceptions and, nevertheless, determines that the incident does fall with the scope of the system then this raises further problems. For example, they must ensure a consistent response to an incident. This is particularly important during the immediate aftermath of an incident when effective action can be taken to mitigate adverse consequences. As we shall see, if these actions are delayed or if 'inappropriate' actions are taken then the net result can be to exacerbate an already serious situation. If the entire decision to investigate an occurrence report is incorrect then this can waster scarce resources and may ultimately convince higher levels of management that the benefits that are derived from the system may not meet the expenditure that is required by the 'false alarms'.

How to Provide Immediate Feedback?

The person who initially receives an incident report must, as mentioned previously, assess its criticality and, if appropriate, must pass it on for further consideration and analysis. If the incident has clear implications for the continued safety of the system then the individual receiving the initial report must directly inform their safety managers so that interim corrective actions can be immediately instigated throughout the organisation. Such notifications have other benefits. While compiling material for this book, I learned of several occasions during which safety managers first learned of critical incidents when a member of the television or press contacted them for their reactions. The notifications of that incident were slowly being passed through the intervening managerment structures of the organisation concerned. Such communications failures have important implications not only for public relations but also for the effective response to incidents and accidents.

Whether or not the immediate recipient of an incident report decides that it falls within the scope of a reporting system, there are two further duties that must be performed in most reporting systems. The first is to inform the contributors, in open or confidential systems, that their reports are being dealt with. This is critical to preserve the trust and coincidence of the participants in the scheme. In the past, completed incident reporting forms have been found in the bottom of supervisors' desks, in pending trays over a month after submission and even in waste paper baskets. One means of avoiding such problems is to develop an auditable paper trail of receipts from the point of submission. This enables those who are responsible for administering a system to trace any potential 'bottle necks'. Many of the organisations, such as the Swedish Air Traffic Control organisation, that exploit these systems have recently turned to electronic implementations that automate the monitoring process and provide staff with feedback on the handling of a report at all stages of the process.

The second documentary obligation on staff receiving an incident report is to provide a written justification of their decision either to proceed with the report, or arguably more importantly, to explain why they decided to drop it. The former is important if incident investigators are to understand why an initial report was passed on for consideration. The later is critical if internal quality control bodies or external regulators are to monitor and approve of decisions that remove potentially critical reports from any subsequent investigation. The importance of this documentation cannot be underestimated. The disclosure laws in several countries make it imperative that such explanations are available. If an initial occurrence report is not investigated and the circumstances of that incident are later replicated by an accident then the potential legal consequences are considerable.

4.1.3 Incident Investigators

Incident investigators conduct the detailed analysis that follows an occurrence report. Rasmussen identified three different diagnostic roles that can be associated with this analysis: analyst, attorney or repairman [696]. These different roles, in part, reflect the difficult of their task. Diagnosis is, however, one one aspect of their duties. They must determine whether any further data acquisition is required, for instance by interviewing more contributors or by examining records from automated logging equipment. Ideally, there investigators work in teams of two or three. This helps to promote the necessary mix of domain-specific, human factors and technical skills. There are also benefits in conducting various interview and elicitation procedures with more than one investigator. The additional expense of forming such groups is, however, beyond the means that are available to many incident reporting systems.

Investigators operate at a local, regional or national level. Given that most reporting systems have a relatively low number of high criticality incidents, many schemes rely upon a small number of highly-skilled investigators. These individuals operate from national or regional centres. However, the additional skills and expertise of such investigators must be balanced against the potential problems of sending 'strangers' to investigate the circumstances of particular incidents in local units. In contrast, other systems have trained larger numbers of investigators who can be appointed from the staff within individual units. This reduces the problems that individuals experience when attempting to understand the working practices of teams that they have not previously met or interacted with. The limitations with this approach are, however, that any investigation can be compromised if the divide between operational and investigatory roles becomes blurred. Investigators may be unwilling to implicate their friends and colleagues. ICAO Annex 13, paragraph 3.1 states that incident investigation is part of the safety improvement process and not part of the operational management of an organisation [384].

What Training do I Need?

The coherent and consistent analysis of occurrences depends upon the careful selection of investigators. They are responsible for drafting the final occurrence report and for submitting it to the appropriate regulatory authority. Recruitment must, therefore, focus on appropriate personality traits. They must be well-organised, meticulous, unbiased, effective communicators etc. These attributes cannot simply be assessed a priori but must be measured and inspected throughout their careers as incident investigators. For example, it is important to determine whether investigators are biased towards certain causal factors in their analysis and interpretation of incidents. It is also important to determine whether investigators continue to consider an appropriate range of recommended remedial actions.

The quality control measures, proposed in the previous paragraph, provide insights into the effectiveness of the training that is provided for incident investigators. Specialised training into the nature and causes of incidents must build upon a detailed knowledge of the working domain. The following list identifies a number of more detailed training requirements:

- *Domain expertise.* Any investigation team must be led by a manager who is competent in the application domain. For instance, it is anticipated that incident investigators will have between five and ten years experience within an Air Traffic Control centre before they are qualified to perform such a role. The meta-level requirement for domain expertise hides a number of more detailed issues. They should understand the working practices of the team that noted the occurrence. They should have a clear view of relevant legislation, regulation and protocols. Their expertise should also be recognised and trusted by employee representatives.

- *Incident investigation expertise.* Chapter 3 has reviewed a number of competing theories and models that describe the ways in which incidents and accidents can occur. The ideas presented in this chapter have different degrees of importance in the training of incident investigators. For instance the previous chapter contrasted Sagan's ideas about high reliability organisations with Perrow's work on normal accidents. It is important for accident investigators to have at least a superficial understanding of these different positions. However, it is essential that accident investigators understand the practical implications of the 'systems approach' to accidents. Similarly, Reason's work on the latent and catalytic causes of failure underpins most recent work in incident investigation.

- *Technical and engineering expertise.* Incident investigators must either posses or have access to specialist knowledge about potential hardware and software failure 'modes'. This is increasingly important as automation enduced failures, typically, emerge from the interaction between a number of component subsystems. It is difficult to under-emphasise the technical challenges that are posed by an investigation and analysis of these incidents. For instance, the integration of application processes can lead to a number of failures that have little superficial connection but which share a number of common causes. Such similarities can only be detected if investigators have considerable technical and engineering skills [413].

- *Human factors expertise.* Given the prominence of human factors in the causation, detection and mitigation of many occurrences, it is necessary to identify a source of human factors expertise for investigators to call upon. This raises a number of pragmatic difficulties. In particular, it is important to emphasise that the analysis of human factors in incident investigation is typically a complex and skilled task. Just as technical and engineering analysis requires competent, specialist training, so does the analysis of human failure. For example, it is often difficult to categorise an error according to a predetermined category. It is critically important to identify and understand those factors that contributed to the error and that helped to shape the operators response to any initial failure.

This is a partial list. More detailed requirements can be identified for particular industries. Additional training requirements can also be identified if investigators must work at the interface between different industries or professions. Air traffic control investigator must understand not the working practices of other controllers but also of pilots. Medical investigators may have to understand the priorities and concerns of several different clinical disciplines.

What Are My Duties?

In order for investigators to complete any analysis of an incident it is important that they have the necessary authority to access all relevant sources of information. This includes immediate access to logs from automated data sources. Investigators must be able to make copies of this information and be able to protect the original logs. They must also have the right to interview key personnel in the aftermath of an adverse occurrence. This can lead to conflict if those members of staff are required for other duties or if they have been excused from duty for psychological or physiological reasons.

Along with these rights, investigators must also fulfill a number of general and specific obligations both to the staff members involved in an incident and to the rest of the safety management structures within their organisation [68]. These general duties include an obligation to ensure a full, independent and objective investigation. To ensure that any investigation and analysis is conducted with the knowledge and participation of operational staff; within the bounds defined by the confidentiality policy that is being used. Investigators must ensure that all relevant documentation is identified, compiled and protected so that subsequent reviews can re-trace the arguments that support their findings. They must also assess the validity and integrity of data that is gathered during any analysis. They must interview all staff who are involved in an occurrenc, again within tbounds specified by the confidentiality policy. They must compile and submit both an initial assessment of the occurrence, typically within 3-10 days of the event, and a final report to their organisation's safety managers. These documents must at a minimum contain an analysis of the occurrence and either interim or final recommendations. They must also ensure that these reports, or a digest, are made available to operational staff so that they can both learn of the outcome of the investigation and see what actions have been identified following from a report.

As mentioned above, investigators must also fulfill a number of obligations that relate more narrowly to the treatment of data that is gathered during any investigation. In particular, access to this data should be restricted to a relatively small number of authorised personnel. If this policy is not enforced then there are strong dangers that data may be lost, corrupted or challenged during any subsequent analysis. It is also important to clearly define permissable uses for the data. For instance, it may only be used for investigating the specific occurrence for which it was gathered. Alternatively, personnel may be told that data will be retained and used to spot emerging trends. In either case, there may be considerable consequences if staff feel that information is retained to monitor individual performance rather than to support more general safety improvements.

4.1.4 Safety Managers

As mentioned previously, each of the roles in this section is generic in the sense that they do not refer to specific posts within a management structure. Instead, they refer to a set of duties or obligations that must be fulfilled in order for an incident reporting system to be effective within an organisation. Safety Managers are ultimately responsible for the operation of the reporting system. They oversee that appointment and working activities of investigators. Together with the regulator, they must also ensure that the recommendations in an occurrence report are acted upon.

How to Resist the Pressure?

Safety managers act as the interface between the investigatory process and many other groups both inside and outside their own organisation. They must propagate information from incident investigators to higher levels of management. This may liaise with training 'departments', with operational staff and with acquisitions groups to ensure that recommendations are implemented throughout the organisation. They must liaise with regulatory authorities and, in more severe incidents, with external investigatory bodies. They may also be expected to liaise either directly with the media or indirectly through public relations organisations. These multiple roles create demands that cannot be underestimated. As noted in previous chapters, they help to account for the way in which local systems are often heavily dependent upon the support of the key individuals that perform this role. In national and regional systems, these pressures can lead to considerable personal stress that may ultimately threaten the success of any incident reporting system. In preparing this book, I interviewed several safety managers who emphasised the invidious nature of their task. They argued cogently that responsibility for the performance of their duties ultimately rested both personally with themselves but also corporately with the directors and managers who must support their actions.

It is very important that safety managers receive adequate protection from the influences that can be exerted on them. For example, it is difficult or impossible to sustain incident reporting functions without a stable budget. This does not imply that infinite resources are required to support the system. It does, however, suggest that frequent cuts without careful planning can and do send inappropriate messages to the staff who must participate in the system. It remains to be seen whether the recent decision to reduce the number of publications of the ASRS' DirectLine journal will have an impact upon the submissions that are made to this system.

Safety managers have further responsibilities. They must protect investigators from undue pressure. External and internal sources can seek to influence the course of an investigation in the hope of having some effect both on the analysis and the recommendations. These pressures can be introduced in covert and discrete ways, through informal meetings, through hints or second-hand reports of the opinions of others within an organisation. In practice, it is difficult or impossible to isolate investigators from these factors. All that safety managers can realistically hope to achieve is to provide investigators with the necessary support so that they can resist the more pernicious influences.

Who Do I Report To?

As mentioned above, safety managers must establish and preserve the communications channels that disseminate lessons from previous incidents. They help to ensure that other groups within the organisation are warned about the potential for similar incidents. This can be done through team briefings, through internal journals or newsletters and increasingly through the electronic media provided by intranets. In practice, however, many of these duties are delegated to incident investigation teams. The safety manager is ultimately responsible for the adequate completion of these tasks.

Safety managers are also responsible for monitoring trend information. For instance, they must encourage participation in an incident reporting system across all geographical regions and managerial groups. They must not only monitor participation rates but must also look for trends of similar incidents that can emerge over time. This task may also involve collaboration with managers of other organisations within the same industry. Of course, this can only be achieved where safety considerations are perceived to be more important than any competitive advantage that might be lost through the exchange of data.

Safety managers must also assess the recommendations that are made by their investigators. Together with operational staff, they can be required to prioritise those recommendations and justify decisions to wither adopt or reject particular findings. They must monitor the implementation of those recommendations that are accepted. They must also monitor the effectiveness of any remedial actions to ensure that they have adequately protected the system against future failures.

Safety managers must prepare briefing documents that are passed to the highest level of management within their organisation. It is, therefore, critical that they have a right of access to upper management. Incident reporting systems are often introduced as a means of improving communication about potential failures within an organisation. The effectiveness of this role will be impaired if all such communication stops with the safety manager. There is also a danger that under such circumstances, managers will only accept recommendations that are amenable to short term fixes [409]. Additional board level support is often required to approve longer term operational changes.

Safety managers must also communicate potential hazards to other groups outside of their own organisations. This can be achieved via a regulatory body. It is important that safety managers have means of communicating directly with the groups or individuals who must intervene to regulate their market. As mentioned in previous chapters, safety managers are, typically, required to provide them with incident statistics. Again, however, safety managers often supplement this information with more pro-active information about wider safety concerns based on their operation experience. A Machiavellian interpretation of this would be that safety managers may predispose the regulator to a positive view of their safety culture. A less cynical interpretation is that this encourages the regulator to fulfill their role as a medium of exchange for safety-related information across an industry.

4.1.5 Regulators

Section 3.1 introduced the role of the regulator by focusing on the ultimate responsibility that they, arguably, hold for failures within an industry. In contrast, this section focuses on the role of the regulator in creating the necessary preconditions for the effective exchange of information through incident reporting. The regulator monitors the performance of the occurrence reporting system as part of the wider safety management processes that are adopted by the management. They often receive copies of all final reports into occurrences as well as reports from the safety managers that describe the measures that have been taken to implement any safety recommendations. Regulators may initiate periodic investigations into particular problems should they continue to receive occurrence reports about similar incidents.

When Do We Regulate?

At a more detailed level, regulators are typically involved in encouraging organisations to establish incident reporting systems. This is often perceived to be part of a wider requirement to encourage safety management programs within their industry. In some sectors, regulators must ensure that organisations meet international obligations:

> "(The assembly) urges all Contracting States to ensure that their aircraft oper ators, providers of air navigation services and equipment, and maintenance organisations have the necessary procedures and policies for voluntary reporting of events that could affect aviation safety" (ICAO Resolution A32-15: ICAO Global Aviation Safety Plan)

However, there are many constraints on the ways in which regulators can intervene to achieve these objectives. As we have seen, OSHA's Cooperative Compliance Programme failed to establish incident reporting as a means of improving safety culture. Employers groups opposed this initiative because it may have placed undue burdens in competitive markets and potentially increased the influence of Federal organisations.

It is again important to emphasise that this section deals with the role and not the office of the regulator. The duties that are associated with regulatory bodies in some industries may, in other industries, be distributed across many organisations. Similarly, they may not be performed at all owing to the nature of the markets that are involved. Healthcare provides an important example of this point. Although some elements of regulation can be associated with the US Food and Drug Administration, there role is primarily focussed on the safety of devices, pharmaceuticals and other products utilised by the medical sector. They do not and have not, typically, been involved in monitoring other adverse occurrences. As a result, the Institute of Medicine report led to the drafting of the Patient Safety and Errors Reduction Act, S.2738, that was introduced into the Senate in 2000. This seeks to establish a national Center for Quality Improvement and Patient Safety under the leadership of a Director who must:

> '(D) develop a confidential national safety database of medical errors reports;
>
> (E) conduct and support research, using the database developed under subparagraph (D), into the causes and potential interventions to decrease the incidence of medical errors and close calls; and
>
> (F) ensure that information contained in the national database developed under subparagraph (D) does not include specific patient, health care provider, or provider of service identifiers.
>
> (2) NATIONAL PATIENT SAFETY DATABASE- The Director shall, in accordance with paragraph (D), establish a confidential national safety database (to be known as the National Patient Safety Database) of reports of medical errors and close calls that can be used only for research to improve the quality and safety of patient care. In developing and managing the National Patient Safety Database, the Director shall–
>
> (A) ensure that the database can only be used for its intended purpose;
>
> (B) ensure that the database is as comprehensive as possible by aggregating data from Federal, State, and private sector patient safety reporting systems;
>
> (C) conduct and support research on the most common medical errors and close calls, their causes, and potential interventions to reduce medical errors and improve the quality and safety of patient care;
>
> (D) report findings made by the Director, based on the data in the database, to clinicians, individuals who manage health care facilities, systems, and plans, patients, and other individuals who can act appropriately to improve patient safety; and
>
> (E) develop a rapid response capacity to provide alerts when specific health care practices pose an imminent threat to patients or health care workers.
>
> (3) CONFIDENTIALITY AND PEER REVIEW PROTECTIONS- Notwithstanding any other provision of law any information (including any data, reports, records, memoranda, analyses, statements, and other communications) developed by or on behalf of a health care provider or provider of services with respect to a medical event, that is contained in the National Patient Safety Database shall be confidential in accordance with section 925.
>
> (4) PATIENT SAFETY REPORTING SYSTEMS- The Director shall identify public and private sector patient safety reporting systems and build scientific knowledge and understanding regarding the most effective–
>
> (A) components of patient safety reporting systems; (B) incentives intended to increase the rate of error reporting; (C) approaches for undertaking root cause analyses; (D) ways to provide feedback to those filing error reports; (E) techniques and tools for collecting, integrating, and analysing patient safety data; and (F) ways to provide meaningful information to patients, consumers, and purchasers'

I view this as a form of regulation because it is an attempt to intervene in the existing market place in a manner that is intended to improve the safety of patients (and staff) within the US healthcare system. In some countries, 'regulation' is a pejorative term that is often associated with ideas of government 'over-regulation'. Those who read the Institute of Medicine report can, however, see that it's authors were careful to balance this fear of intrusion in the marketplace against the need to address the consequences of human error in medicine [481]. Those same concerns are apparent in this draft of the Act.

Not only must regulators help to establish incident reporting systems, they must also monitor their operation. As we shall see, this is a non-trivial exercise. There is the obvious paradox that a relatively low number of reported incidents may indicate a high degree of safety within an organisation or a relatively low participation rate. Similarly, it can be difficult to determine whether the investigatory procedures that lead to a criticality assessment of each incident are implemented in the same manner across different organisations. For instance, some European Air Traffic Service provides classify the severity of an incident according to its worst plausible outcome . An air proximity violation that was resolved by the actions of the crew might, therefore, be treated as if a collision had occurred because ATS personnel had not intervened to avoid the incident from become more serious. Other organisations within the same industry would treat this as a far less serious incident. Under this vew, the aircrew are perceived to form part of the wider safety system. A collision was avoided and hence that system functioned as intended.

Regulators must intervene to support the exchange of safety-related information throughout an industry. This responsibility is a repeated theme in the Patient Safety and Errors Reduction Act cited above. However, it can also be seen in the regulatory structures that govern other industries. For instance, the regulatory safety functions of the UK rail industry are performed by the Railways Inspectorate within the Health and Safety Commission of the Health and Safety Executive. In contrast, the economic functions associated with performance measurement, standard setting and price monitoring are performed by the office of the Rail Regulator. The regulatory role of the Health and Safety Commission in establishing confidential incident reporting schemes can be seen in their action plan to implement the recommendations of the recent inquiry into the Southall rail crash:

> "All parties in the rail industry should co-operate in the collection of evidence to support reliable research into human behaviour studies relating to driver performance. Railtrack should co-ordinate this work and TOCs (Train Operating Companies) incorporate the results into training programmes (paras 1.25, 7.16, 16.2).
>
> Evidence should include that to be provided by CIRAS (Confidential Incident and Reporting System) and from On-Train Data Recorders used to monitor driver behaviour. ASLEF (Associated Society of Locomotive Engineers and Firemen) in particular should give their full support to such an initiative (paras 14.23, 14.25, 15.15, 16.3).
>
> Comment: Much of the information required for the human factors work on driver behaviour (Recommendation 1) will be provided by train operators. Most TOCs have already agreed to enroll their drivers in CIRAS (or equivalent confidential reporting systems) following the Rail Summit; coverage should be complete nationally with all staff briefed by 1 April 2001. Individual TOCs agree to interrogate and provide data analysis of on-train data recorders or from other available means of recording driver activity. ASLEF and RMT (National Union of Rail, Maritime and Transport workers) support approach, subject to confidentiality reassurances. HSC (Health and Safety Commission) agrees that this action should be on individual train (both passenger and freight) operators. Action: Individual TOCs to submit a progress report to HSC confirming their active participation in providing human factors data to Railtrack and enrollment of driver in CIRAS. ATOC (Association of Train Operating Companies) to set up a system to identify good practice on how driver behaviour is to be monitored using OTDRs (On train data recorders). Progress report to be submitted to the HSC." [317].

The Health and Safety Commission are intervening to ensure that all parties in the rail industry cooperate to collect evidence about the human factors problems that affect driver performance. All train operating companies must establish a confidential incident reporting system, similar to CIRAS mentioned in previous chapters. These companies must, in turn, agree to provide access to the data that is obtained by these systems.

The previous paragraph has argued that regulators play a role in the collection and dissemination of information within an industry. This may clearly involve a delicate balance between the promotion of safety and the exchange of commercially sensitive information. This balancing act becomes even more complex when regulators attempt to promote the exchange of information across national boundaries:

> "The Board's focus extends beyond the United States' borders. Realizing that chemical accidents may have global health, environmental and economic effects, Congress encouraged the Board to offer investigative assistance to other countries, both as a means of helping and as a method of learning. Through its international outreach efforts to

> government and industry, the Board can ensure its safety research program, professional services and technical information accurately and adequately address the world's chemical safety needs". [163]

The sensitivity of the information that is often provided by incident reporting systems perhaps accounts for the notable lack of success in achieving the international collaboration that many regulators envisage. However, this view is being challenged by recent commercial initiatives to encourage the exchange of occurrence data within the aviation industry [308]. The GAIN system, introduced in Chapter 2 has over the last three years been transferred away from the FAA to the airline industry itself [680]. At present, GAIN simply acts as a clearing house for data gathered by other public sources including the ASRS and FAA incident reporting schemes. In the future, however, it may provide greater opportunities for the exchange of data directly between aviation operating companies even though that data is unlikely to be publically accessible to the same extent as the ASRS sources.

What Information Do We Need?

The previous section described the role of the regulator in setting up and monitoring incident reporting systems. They must also ensure that the output from these systems is collated, analysed and effectively used to address safety-related problems that arise across an industry or between several industries. This duty can be stated relatively simply. However, it is far harder to achieve. This difficulty of performing this task can be illustrated by the FDA's Manufacturer and User Facility Device Experience Database (MAUDE) [272]. This tool represents a significant advance on many existing regulatory incident reporting systems because it provides manufacturers and operators with an accessible means of looking for information about previous incidents. Techncial details about this system will be provided in Chapter 14 and the interface to this system is illustrated in Figure 14.4. Users can access incident data in the MAUDE database by selecting a number of predefined categories or by entering a free-text search. The following quotation illustrates the types of data that can be retrieved using this system:

> "Adverse event or product problem description: A susceptibility report message, which the microbiology lab uses as an indicator for verifying oxacillin results did not print on a pt lab report from the ... instrument. The lab did not verify ... results on this pt's blood culture report. The pt's physician has stated that as a result of a lab error, a treatment error occurred leading to development of an abscess. This abscess has put pressure on the pt's spinal cord causing paralysis of the legs.
>
> Additional manufacturer narrative: An investigation into the customer complaint determined that a message, which previously printed on the instrument lab report, no longer prints with the release of a new software version.
>
> ...[It could not be concluded] that lack of this message caused or contributed any negative effects to the pts condition ased on the following points: 1. subsequent blood cultures were negative after treatment with oxacillin. 2. the lab did not save the original isolate from the blood culture used for testing on the system. 3. this pt had a previous history of a oxacillin resistant staph aureus infection. 4. treatment of an abscess, regardless of culture and susceptibility results, routinely requires more intervention than simply administering antibiotics. The message for oxacillin will be added to the lab reports with the next software release. All customers will be notified immediately by letter concerning the missing message when the oxacillin indicator antibiotics are resistant"

The MAUDE system is important because it illustrates the ways in which regulators can intervene to act as a clearing-house for incident data. The development of search engines, for the first time, provides users with the opportunity to identify common trends across an industry. However, the previous examples also illustrate some of the challenges that are facing such regulatory action. In particular, the previous search for software related incidents yielded five hundred hits amongst the MAUDE collection. At this point the system halted its search and prompted me to refine my search because there was too much relevant data. In some senses this reflects the way in which incident reporting systems can become victims of their own success. For instance, the ASRS system know holds over 500,000 records. Later section will describe software engineering and information management tools that can be used to address these problems and still enable users to identify common trends amongst a growing mass of incident data.

4.2 Different Anatomies

The previous section has summarised a number of the key roles that support incident reporting. In contrast, however, this section goes beyond these roles to look at a number of different reporting architectures. These architectures reflect the organisation that is necessary to collect incident reports, analyse them and then make recommendations. Clearly, the managerial structures that are necessary to support large national and international systems are unlikely to be appropriate or even necessary in smaller scale local and regional systems. This section, therefore, provides a brief overview of a number of different ways in which incident reporting systems can be managed.

4.2.1 Simple Monitoring Architectures

Figure 4.1 represents the simplest architecture for an incident reporting system. A contributor submits a report based on the occurrence that they have witnessed or are concerned about. This submission process can be implemented using printed forms, by telephone calls, or increasingly using computer-based techniques. An external agency received the report and after assessing whether or not it falls within the scope of the system they will decide whether or not to publish information about the occurrence. The contributor and others with the same industry can then read the report and any related analysis before taking appropriate corrective actions.

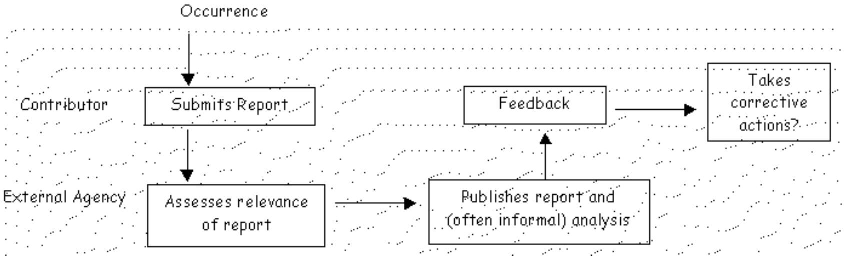
Occurrence
Contributor
Submits Report
Feedback
Takes corrective actions?
External Agency
Assesses relevance of report
Publishes report and (often informal) analysis

Figure 4.1: A Simple Monitoring Architecture

This approach is typified by the Swiss Confidential Incident Reporting in Anaesthesia system (CIRS) [755]. A web-based form is used to submit an incident report to the managers of the system. Given the sensitive nature of these incidents, this is an anonymous scheme. The managers cannot, therefore, conduct follow-up investigations. However, they do perform a high-level analysis of this and similar events before publishing a summary on their web site.

There are a number of limitations with the architecture shown in Figure 4.1. In particular, this simple monitoring approach simply provides a means of disseminating information about previous failures. There are no guarantees that individual organisations will take any necessary corrective actions. Similarly, there is a danger that different institutions will respond to the same incident in different ways. This inconsistency creates the opportunity for future failures if an organisation fails to correctly safeguard the system. A further problem is that this approach does not provide any means of determining whether reports were accurate or not. This creates potential dangers because a report may omit necessary information about the causes of an incident. As a result, other organisations might respond to the symptoms rather than the underlying problems that lead to an occurrence. As most of these systems are truly anonymous, it can be difficult or impossible for the managers of the scheme to identify whether any local, contextual factors contributed to an incident.

As with all of the architectures presented in this section, there exist a number of variations that have been used to structure existing systems. For instance, the US Food and Drug Administration's MAUDE system, mentioned above, cuts out the external agency and enables individuals to report directly to the regulator. These reports are then posted on the FDA's web site. If the incident is considered serious enough then the regulator may intervene through a product recall or amendment notice.

4.2.2 Regulated Monitoring Architectures

Figure 4.2 provides a high-level view of what we have termed the 'regulated monitoring' architecture for incident reporting. This is very similar to the approach described in the previous section. However, in this approach the external agency that received the contribution can go back and ask further questions to refine their understanding of an occurrence. Once they are clear about what has taken place, they produce a summary report that, typically, does not reveal the identity of their contributor. This summary is then placed before management and regulators who are responsible for identifying corrective actions. They must also determine whether those corrective actions can be implemented. The reporting agency will then receive a report on corrective actions that can then be communicated back to the original contributors and their colleagues through journal or newletter publications.

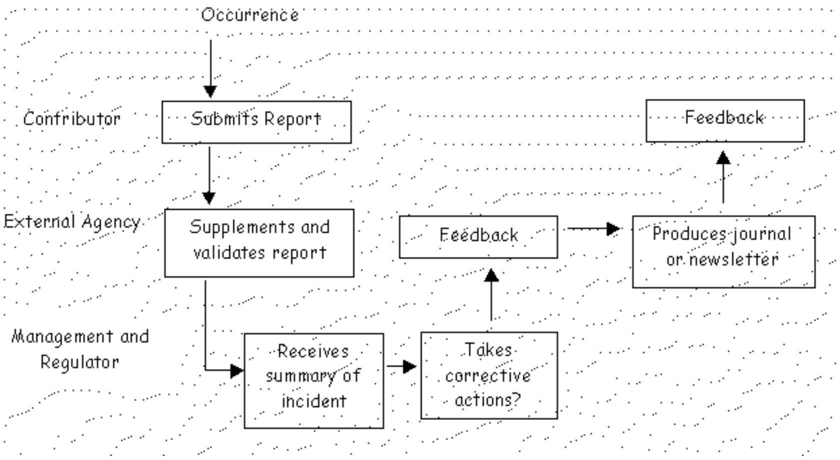
Occurrence
Contributor
Submits Report
Feedback
External Agency
Supplements and validates report
Feedback
Produces journal or newsletter
Management and Regulator
Receives summary of incident
Takes corrective actions?

Figure 4.2: Regulated Monitoring Reporting System

The Confidential Incident Reporting and Analysis System (CIRAS) is a good example of an incident reporting scheme that implements the high-level architecture illustrated in Figure 4.2. This receives paper-based forms from Scottish train drivers, maintenance engineers and other rail staff. A limited number of personnel are responsible for processing these forms. They will conduct follow-up interviews in-person or over the telephone. These calls are not made to the contributor's workplace for obvious reasons. The original report form is then returned to the employee. No copies are made. CIRAS staff type-up a record of the incident and conduct a preliminary analysis. However, all identifying information is removed from the report before it is submitted for further analysis. From this point it is impossible to link a particular report to a particular employee. The records are held on a non-networked and 'protected' database. This data itself is not revealed to industry management. However, summary reports are provided to management at three monthly intervals. This concern to preserve trust and protect confidentiality is emphasised by the fact that a unit within Strathclyde University employs the personnel who process the reports rather than the rail operators.

The FAA's ASRS provides a further example of the architecture illustrated in Figure 4.2. NASA plays the role of the external reporting agency. Feedback is provided through a number of publications, such as the Callback newsletter and the DirectLine journal. An important strength of the publications produced by this approach is that it provides a measures assessment of several incidents through the editors' analysis. It also enables staff to read an explanation of an incident through the words of their colleagues.

Again there are a number of limitation with the high-level architecture shown in Figure 4.2. These do not stem principally from the problems of accessing more detailed causal information, as was the case with simple monitoring architectures. In contrast, they stem from the additional costs and complexities that are introduced by external reporting agencies. In particular, it can be difficult to preserve an independent but co-operative relationship between the organisation's management and a reporting agency. This relationship can become particularly strained when the agency is responsible for identifying corrective or remedial actions that the management must then implement. The ALARP (as low as reasonably practicable) principle is often used to justify resource allocation. The subjective nature of this approach can lead to conflicts over the priority allocated to many remedial actions. There is also a danger that these schemes will resort to low-cost reminders [409]. In consequence, many schemes operate on a smaller-scale, more local level. These schemes rely upon the same individuals to both collect the data and take immediate remedial actions.

4.2.3 Local Oversight Architectures

Figure4.3 illustrates the architecture that typifies many locally operated, incident-reporting systems. In many ways, these schemes were the pioneers of the larger more elaborate systems that have been mentioned in the previous sections. Individual sponsors either witness other schemes or independently decide to set up their own. Staff are encouraged to pass on incident reports to them. Typically, this is in a confidential rather than an anonymous fashion. Even if the forms do not ask for identification information it is often possible for the sponsors to infer who is likely to have submitted a form given their local knowledge of shift patterns and working activities. The sponsors can supplement the reports from their own knowledge of the procedures and practices within a unit. This enables them to analyse and validate the submission before passing a summary to their management. In contrast to other architectures, however, they are in a position to take direct remedial action. This is, typically, published in a newsletter. These publications not only provide feedback, they are also intended to encourage further submissions.

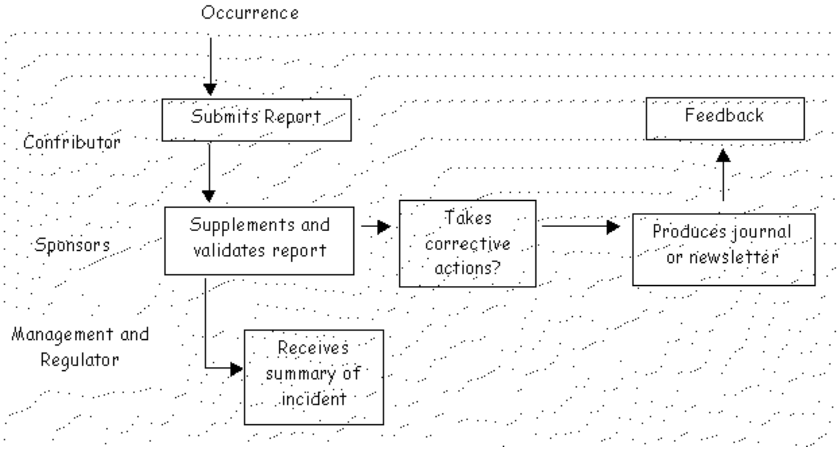
Occurrence
Contributor
Submits Report
Feedback
Sponsors
Supplements and validates report
Takes corrective actions?
Produces journal or newsletter
Management and Regulator
Receives summary of incident

Figure 4.3: Local Oversight Reporting System

Local oversight architectures are illustrated by one of the longest running medical incident reporting systems. David Wright, a consultant within the Intensive Care Unit of an Edinburgh hospital, established this system over ten years ago [119]. The unit has eight beds at its disposal with approximately three medical staff, one consultant, and up to eight nurses per shift on the ward. David Wright receives each report. They are then analysed with the help of a senior nurse. Any necessary corrective actions are instigated by them. Trust in the sponsor of this system is a primary concern, given the relatively close-knit working environment of an intensive care unit. The success of the system depends upon their reputation and enthusiasm. The extent of his role is indicated by the fact that less reports are submitted when David Wright is not personally running the system. The reports from these systems provide a valuable insight into problems in the particular practices and procedures that are followed within an organisation.

The strengths and weaknesses of such local systems are readily apparent. The intimate local knowledge and direct involvement with the contributors makes the interpretation and analysis of incident reports far easier than in other systems. However, it can be difficult to replace key personnel and sustain confidence in the system. It can also be difficult to drive through deeper structural or managerial changes from local systems. Individual sponsors often lack the necessary authority (or resources) to instigate such responses. As a result they often 'target the doable'. Similarly, it can be difficult to co-ordinate the efficient exchange of date between local systems to get a clearer overview of regional, national and even international trends.

4.2.4 Gatekeeper Architecture

Figure4.4 illustrates the architecture of several national incident-reporting systems. The increased scale of such systems usually implies the greater degree of managerial complexity apparent in this framework. The contributor submits a report to their local manager. They may then take some initial remedial actions and then passes the form to a 'gatekeeper'. They register the report; in any national system there is a danger that individual contributions may be lost or delayed. The gatekeeper has this name because they must determine whether the occurrence is important enough to allocate further analytical and investigatory resources. If this decision is made then they will delegate the report to another unit within the organisation that is responsible for the aspect of the system that was most directly affected by the occurrence. The report is passed to a handler within this service department and they attempt to identify means of resolving any potential problems. Feedback is then provided to the contributor via their local manager. This approach is, typically, confidential or open rather than anonymous.

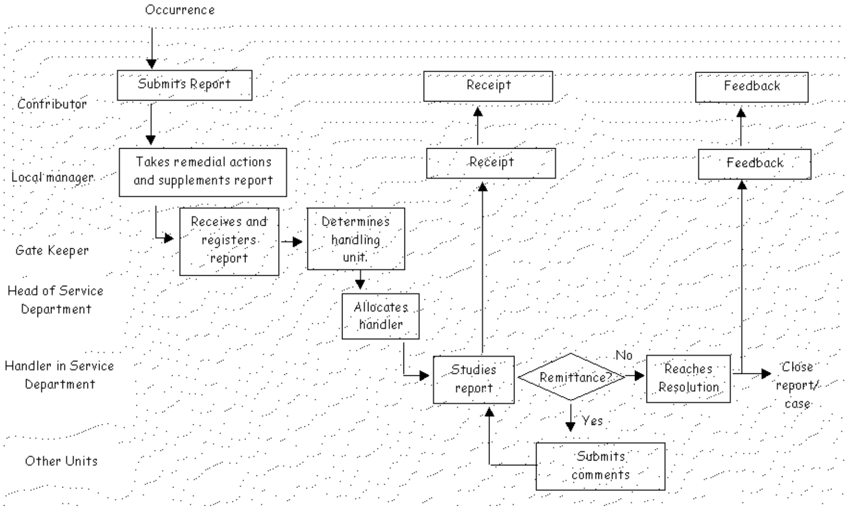

Figure 4.4: Gatekeeper Reporting System

This approach is exploited by the Swedish Air Traffic Control system. It is unusual in that it encourages the open reporting of a wide range of potential and observed failures. The definition of an 'occurrence' includes all forms of human, operational and technical failures even including incidents such as a failure of a light bulb. All reports are handled centrally by a number of specially trained gatekeepers who are responsible for filtering the reports and then passing them on to the relevant departments for action.

These individuals must be highly trained both in the application domain of air traffic control but also in the technical problems that lead to system failures. However, because all occurrence reports pass through their offices they gain a detailed understanding of both operator behaviour and system performance. The gatekeepers, therefore, are in a position to provide valuable information both to training directors but also to the risk assessments that guide future investment decisions.

The gatekeepers are an important strength of the system shown in Figure4.4. They are responsible for filtering reports and allocating remedial actions. This centralisation ensures a consistent analysis and response. However, they are a critical resource. There is a risk that they may act as a bottleneck if incidents are not handled promptly. This is particularly important because delays can occur while reports are sent from outlying areas to the gatekeeper's central offices. The Swedish system has addressed many of these criticisms by adopting a range of computer-based systems that keep safety managers and contributors constantly informed about the progress of every incident report. However, there remains the danger for many of these systems that any omissions in the training of a gatekeeper can result in incorrect decisions being made consistently at a national level.

4.2.5 Devolved Architecture

Figure4.5 provides an overview of an alternative architecture for a national system. Rather than have a central gatekeeper who decides whether an incident falls within the scope of the system, this approach relies upon a more decentralised policy. Any of the personnel involved in the system can decide to suspend an investigation providing that they justify their decision in writing and pass their analysis to the safety management group who monitors the scheme. As can be seen, contributors pass their reports to their supervisors. This is important because in many industries, such as air traffic control, the individuals who are involved in an incident will often be relieved of their duties. A sense of guilt can often affect their subsequent performance and this can endanger further lives. In national systems, it is often common to provide an alternative submission route through an independent agency in case a report is critical of the actions taken by a supervisor.

The supervisor takes any immediate actions that are necessary to safeguard the system and informs the safety management group if the incident is sufficiently serious. The safety management group may then commission an initial report from a specialist investigation unit. They may also decide to provide an immediate notification to other personnel about a potential problem under investigation. These investigators may call upon external experts. Depending on the conclusions of this initial report they may also be requested to produce a final report that will be communicated back to the safety management group. In a number of these systems, final reports are issued to the original contributors who can append any points of further clarification. The safety management group is then responsible for communicating the findings and for implementing any recommendations following discussions with the regulatory authorities.

Figure4.5 illustrates the complexities involved in organising nation and international reporting systems. It depends upon the co-ordination and co-operation of many different individuals and groups. However, such architectures are necessary when the problems of scale threaten to overwhelm systems based on the approach illustrated in Figure4.4. The problem with this system is that there is a greater chance of inconsistency because different staff determine how an occurrence is to be reported and investigated. Different supervisors may have different criteria for what constitutes an occurrence that should be passed on for further investigation. Most European air traffic control service providers have tackled this problem by publishing exhaustive guidelines on what should be reported. These guidelines are distributed to all personnel and are addressed during the training of control staff.

It is important to emphasise that this section has avoided normative arguments about the absolute value of the different architectures that have been presented. This is entirely deliberate. As suggested in the previous paragraph, we know very little about the impact of these different man-

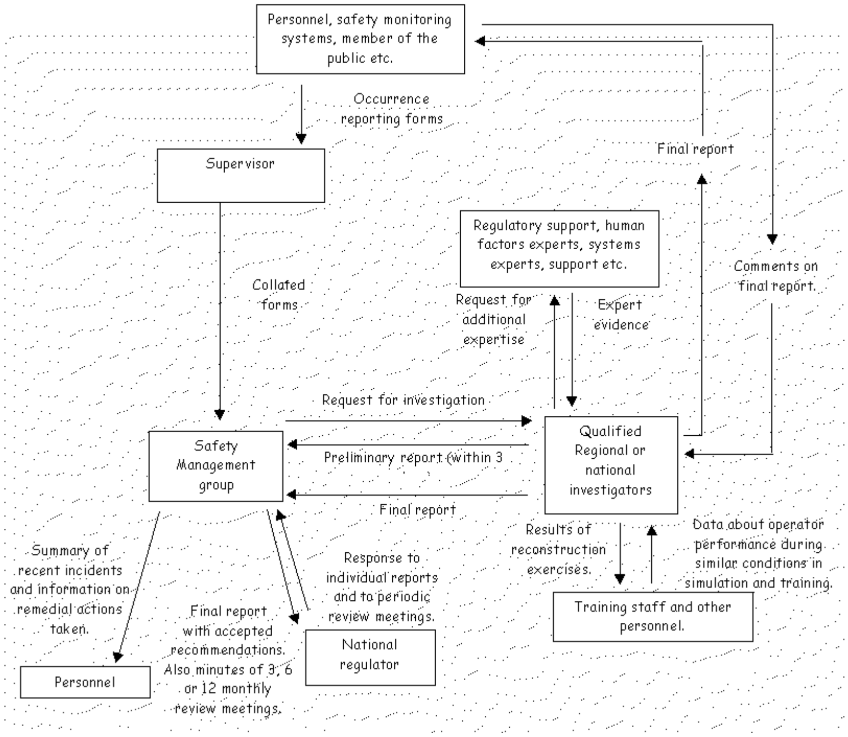
Personnel, safety monitoring systems, member of the public etc.
Occurrence reporting forms
Supervisor
Final report
Regulatory support, human factors experts, systems experts, support etc.
Comments on final report.
Collated forms
Request for additional expertise
Expert evidence
Request for investigation
Safety Management group
Preliminary report (within 3
Qualified Regional or national investigators
Final report
Results of reconstruction exercises.
Data about operator performance during similar conditions in simulation and training.
Summary of recent incidents and information on remedial actions taken.
Response to individual reports and to periodic review meetings.
Training staff and other personnel.
Final report with accepted recommendations. Also minutes of 3, 6 or 12 monthly review meetings.
National regulator
Personnel

Figure 4.5: Devolved Reporting System

agement structures. In consequence, it is difficult to be confident in any comparative analysis. Tools and techniques for performing such comparisons are urgently needed as incident reporting systems continue to proliferate in many different industries.

4.3 Summary

This chapter began by considering a number of different roles that together contribute to the successful operation of many incident reporting systems. These roles are generic in the sense that they represent key activities during the reporting, analysis and subsequent implementation of safety recommendations. These activities may be associated with particular individuals or with teams depending upon the scale and the organisation of the reporting scheme.

This chapter initially focussed on the reporting of adverse occurrences by individuals and groups in the workforce. The opening sections focussed on the rights and duties of these contributors. The next chapter will provide a greater consideration of the ways in which automated monitoring equipment can be used in the detection of adverse occurrences.

The following sections went on to consider the triage that is required when a report is initially received. Line managers and supervisors are, typically, required to secure the short term safety of the system in the aftermath of an incident. They must also pass on reports so that they can be processed in a prompt and efficient manner.

This chapter also looked at how information must be passed to incident investigators. We considered the powers that investigators must have if they are to elicit relevant information about an occurrence. Later sections also considered the training requirements and professional obligations that must be met by these individuals.

We went on to consider what we have described as the ‘invidious’ role of the safety manager. They act as a conduit of safety information from the workforce to higher management. They must also effectively communicate safety objectives from higher levels of management down to the workforce.

It was stressed that they must communicate effectively not only within their organisation but also to external bodies including industry regulators.

Regulators have been defined as organisations that intervene in the normal operation of the market to achieve economic and social objectives, such as improved safety, that might otherwise be overlooked. This chapter examined the tensions that arise when regulatory actions must balance the need to exchange safety information against the danger of forcing companies to pass on what might be commercially sensitive data. We also briefly considered nascent attempts by a commercial consortium to encourage the global exchange of incident information.

The second half of this chapter then went on to look at how these different roles contributed to different types of incident reporting system. Simple monitoring architectures simply provide a common point of access to incident reports. They are, typically, anonymous and so only a cursory validation can be performed. There are a number of limitations that restrict the utility of these systems, although they are simple to operate and can be established at low-cost. In particular, there is no guarantee that the submissions are genuine nor is there typically any guarantee that different institutions or investigators will arrive at a consistent interpretation of the events that are described.

Regulated monitoring architectures extend simple monitoring architectures by introducing an external agency that intervenes to validate and supplement any initial report. This additional validation increases the range of evidence that is available within the system and also helps to support any subsequent analysis of an adverse occurrence. However, the costs of maintaining such an external investigatory body are typically beyond the resources of most local systems.

Local oversight architectures rely upon key individuals or sponsors who can use their knowledge of a working environment to interpret and assess the reports that they receive. These individuals may perform additional validation and investigation but this need not always be necessary depending in their involvement in the target system. However, there is a danger that such systems are susceptible to the personal biases and training of these key sponsors. It can also be difficult to reestablish staff trust in any system when these individuals leave or take up other duties.

Gatekeeper architectures represent a more complex architecture in which a number of key individuals together perform the triage that might otherwise have been performed by a single individual within a local system. These individuals are trained to identify the severity of a report and to allocate a handling unit that is tasked to respond to that incident. However, in national and regional systems there is a danger that they may act as a bottleneck if incidents are not handled promptly. This is particularly important because delays can occur while reports are sent from outlying areas to the gatekeeper's central offices.

Finally, devolved architectures are intended to support large-scale national and international incident reporting systems. Information is passed through the different levels of an organisation up to a body of 'professional' incident investigators. These investigators report to the safety managers, mentioned above. Elaborations on this approach include several feedback mechanisms so that contributors are continually involved in the investigation and analysis process. Again, however, the costs associated with such a system would dissuade many industries from adopting every aspect of this approach.

As mentioned, the intention in this chapter has not been to recommend any particular architectures. In contrast, the intention has been, for the first time, to provide an overview of the different approaches that are currently being used within individual reporting systems. The following chapter builds on this analysis and begins to look in detail at key stages in the operation of an incident reporting system. These include: detection and notification; data gathering; reconstruction; analysis; recommendation and monitoring; reporting and exchange. As before, the intention is to provide a generic analysis of activities that are common to many different types of system. It is also intended the analysis provide pragmatic advice and guidance based on comparative studies of systems in several different industries.

Chapter 5

Detection and Notification

The previous chapter presented a number of different ways in which incident reporting systems can be organised. These architectures ranged from small-scale local systems through intermediate gatekeeper systems through to more complex, devolved, national and international mechanisms. The following chapters build on this by examining a number of generic problems that must be addressed by all incident reporting systems. These issues are illustrated in Figure 5.1. As can be seen, the detection and notification of an occurrence is followed by a phase in which data is gathered about the events leading to a failure. This data can be used to reconstruct the likely ways in which events combined during the course of an incident. Once a probable reconstruction has been developed, it is possible to analyse these likely scenarios to identify key latent and catalytic causes. These form the focus for any subsequent recommendations about ways to prevent future failures. If these recommendations are adopted then they must be acted upon and their outcomes must be monitoried.

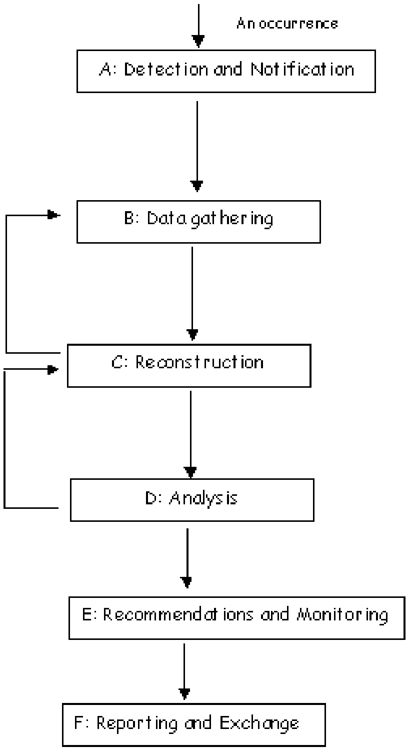
An occurrence
A: Detection and Notification
B: Data gathering
C: Reconstruction
D: Analysis
E: Recommendations and Monitoring
F: Reporting and Exchange

Figure 5.1: Generic Phases in Incident Reporting Systems

Clearly, it is important to determine whether any potential improvements are actually delivering the anticipated benefits. Finally, information about incidents must be reported to others both inside and outside an organisation.

Figure 5.1 includes two lines of feedback. Once investigators begin a period of reconstruction, they may often identify the need for further information about the course of an incident. In other words, they may be forced to continue with data gathering exercises. For example, it may not be possible to immediately determine what key individuals or systems were doing during particular stages of an occurrence. Investigators must, therefore, go back and conduct further interviews or extract additional system logs where they are available. Similarly, the analysis of an occurrence can often help to identify inconsistencies or omissions in the reconstruction of an occurrence. Assumptions about the flow of events leading to a failure may be proved incorrect or implausible during the later stages of an investigation.

As in previous chapters, Figure 5.1 makes no assumptions about the managerial structures that are used to implement these phases. For example, in a national confidential system the data gathering phase may consist of trained field investigators calling on a working group to interview members of staff who were involved in an occurrence. In a small-scale anonymous system, data gathering may involve less formal conversations with personnel in similar working environments to determine whether the concerns in the occurrence report are shared by the other colleagues. Clearly, the sophistication, organisation and investment involved in each of the stages also depends upon the scale of the reporting system. As we shall see, national and international schemes may deploy sophisticated three-dimensional, immersive virtual reality simulators to reconstruct the events leading to particular failures. Such an approach is, typically, beyond the resources of most local systems.

5.1 'Incident Starvation' and the Problems of Under-Reporting

This chapter begins our analysis of the generic phases shown in Figure 5.1 by focusing on the problems of detection and notification. Some of the concerns that arise during this initial stage are illustrated by the UK's guidelines for reporting adverse incidents with medical devices:

> "All staff, including contractors, should be regularly reminded of their responsibilities with regard to adverse incident reporting and of the relevant local procedures including the need to isolate and retain defective or suspect items. This information should also be conveyed to new staff as part of their induction training. The procedures should ensure that: where appropriate, a liaison officer is appointed with the necessary authority to take responsibility for the reporting of medical device related adverse incidents to the Medical Devices Agency (MDA) as detailed in the Annexes; devices involved in an adverse incident together with other material evidence (e.g., packaging of a single use device) should be clearly identified and kept in quarantine, where practicable, until MDA's device specialists have been consulted. Where quarantine is not practicable, the state of the device(s) at the time of the incident should be recorded for use in any subsequent investigation; local action is taken as necessary to ensure the safety of patients, users and o thers. Regular reviews should be undertaken to ensure that the procedures are effective and are being followed." [535]

As this quotation suggests, workers must receive training about what to report and how to report it. Setting up the necessary infrastructure for an incident reporting system does not guarantee that staff will be motivated to participate. This excerpt also stresses the importance of local liason officers, even in a large national reporting system. These trusted advocates support staff who are concerned about adverse occurrences. They must address contributors' concerns about anonymity and confidentiality that were described in Chapter 4 as part of a more general review of the key roles that support incident reporting systems. The net effect of these concerns is to exacerbate problems of under-reporting. Rather than reiterate the importance of addressing contributors' concerns about anonymity and confidentiality, the following paragraphs look at techniques that are specifically intended encourage the notification of adverse occurrences.

The ultimate aim of incident reporting systems is to identify the causes of previous failures and to use this understanding to avoid or reduce future problems. Demonstrating such 'improvements' is complicated because voluntary incident reporting systems often suffer from *chronic* under-reporting. The fear of retribution and the concern that reports will not be acted upon have dissuaded individuals from contributing to a system. The reality of incident reporting in the UK NHS is illustrated by the report into the Royal College of Anaesthetist's critical incident system:

> "We know from previous studies that self-reporting of incidents retrieves only about 30% of the incidents reported by independent observer. We do not know, therefore, either true numerators nor because we do not collect them, denominators; even the Department of Health does not know how many anaesthetics are given annually. Any idea that this scheme might give absolute incident rates must therefore unfortunately be rejected. what we can hope to do is to paint a picture of what we are told nationally and allow departments to see whether the incidents that they are seeing locally are common or rare..." [715]

Vincent, Taylor-Adams and Stanhope observe that between 4-17% of patients in acute hospitals studies suffer from iatrogenic injury [849]. Observational studies have found that 45% of patients experienced some medical mismanagement and 17% suffered events that led to a longer hospital stay or more serious problems [28]. It has been estimated that approximately 850,000 adverse events occur within the UK National Health Service each year [633]. The earlier Harvard Medical Practice Study used similar techniques to estimate that among the 2,671,863 patients discharged from New York hospitals in 1984 there were 98,609 adverse events and 27,179 adverse events involving negligence [93]. Even the most successful voluntary reporting systems only succeed in eliciting information about a tiny fraction of the incidents that are revealed by the exhaustive analysis of records and logs. For instance, Barach and Small estimate that between 50 and 95% of medical incidents go unreported [66].

5.1.1 Reporting Bias

To summarise, targets for the reduction of incidents, such as those proposed by the UK NHS, depend upon a bench-mark assessment of existing incident rates. Incident reporting systems provide useful information about the causes of some incident. However, they do not provide accurate assessments of background frequencies. Alternative techniques must be used to calculate these incident rates. These can be summarised as follows:

1. *extrapolation based on snap-shot samples.* The key technique that drives most base-line estimates of incident frequency is to extrapolate from exhaustive analysis of small samples. This approach, however, is fraught with analytical problems. Clearly, the sample size and selection is a critical issue. If these are in any way biased or unrepresentative then the results of any analysis will be flawed. Further problems stem from the sorts of data that comprise such a sample. There are few guarantees that logs and records will provide indications of all potential incidents. If they do not then a further source of under-reporting is introduced. If observational techniques are used, in which analysts directly monitor work tasks, then there is a danger that the presence of the analyst will itself distort normal working practices;

2. *post hoc analysis of logs and other data recordings.* Exhaustive searches can be made through all of the data that may have been amassed during a specified operating interval. This information can be manually assessed to determine whether or not it provides evidence of a potential adverse incident. Although this might seem to be a relatively straightforward task, there are numerous complications. In air traffic control, the physical separation between aircraft can be calculated from radar logs. However, this would be impracticable in the general case given the volume of aircraft movements in most sectors. Such an analysis would not also indicate errors of intention or lapses that were rectified before an infringement actually occurred. Similar problems arise in the medical domain. Inadequate and partial record keeping can make it

difficult to determine whether or not an error was actually made or if that error actually had any observable clinical consequences;

3. *automated incident detection* Clearly, the burdens of manually search for indications of incidents can frustrate attempts to obtain clear base-line measures of incident rates. As a result, a variety of automated tools (see below) can be used to search for key indicators. These tools range from simple databases through to more advanced data mining systems similar to those that will be discussed in Chapter 15. However, such tools introduce a further level of indirection that can bias results in ways that are often difficult to predict. In particular, there are the twin problems of precision and recall. A low precision search will detect many potential incidents that analysts must manually assess and then reject as not representing actual incidents. A low recall search will yield a number of potential incidents but will also leave many real cases undetected in the mass of incident data.

4. *observational studies.* Finally, as mentioned above, observational studies can be used to identify background statistics for the numbers of adverse occurrences within an organisation. This relies upon trained analysts monitoring everyday activities to detect adverse occurrences 'on the job'. This approach has yielded many important insights into other areas of human-system interaction. However, there are considerable practical problems in applying it to assess incident frequencies. Previous paragraphs argued that workers will adjust their behaviour if they believe that they are being monitored. This has been termed the Hawthorne effect after the 1939 study of workers in the Western Electric Company's plant in Hawthorne Illinois. Productivity rose shortly after investigators started to observe workers even before any changes were made to working patterns. Other problems relate to the limited scope and high costs that can be associated with observations techniques. In particular, the low frequency of some types of incidents may mean that a team might have to continue to observe activities for many months before an incident is detected.

Jha, Kuperman, Teich, Leape, Shea, Rittenberg, Burdick, Segerand, Vander Vliet and Bates hae conducted several studies into the use of both manual and automated techniques for assessing base-line incident frequencies [400]. Most of their work focuses on adverse drug events which they argue are both common and costly. They criticise the 'spontaneous', voluntary systems in most hospitals as lacking sensitivity. They also criticise the costs associated with the exhaustive manual analysis of patient charts. As a result, they have worked to develop a computer-based adverse drug event monitor. Subsequent studies have then compared the performance of this tool with the products of both chart review and voluntary report systems. In one study, they focused on all patients admitted to nine medical and surgical units in an eight-month period [400]. The monitoring program identified situations that suggested a potential adverse drug event. These included requests for antidotes, such as naloxone. A trained reviewer then examined the patient's records to determine whether an adverse incident had occurred. The results were then compared with the products of an intensive manual review and a voluntary reporting system operated by nurses and pharmacists. Both the automated system and the chart review strategies were independent, and the reviewers were blinded.

The computer monitoring strategy identified 2,620 of which only 275 were determined to be adverse drug events. This illustrates the problems of poor precision, mentioned above. The manual review found 398 adverse drug events, whereas voluntary report only detected 23. Of the 617 ADEs detected by at least one method, manual review detected 65%, the automated program identified 45% and voluntary reporting contributed only 4%. It can be argued that all three techniques suffered from the problems associated with poor recall. This work has clear and profound implications for managers and regulators who must encourage participation in incident reporting systems:

> "The computer-based monitor identified fewer Adverse Drug Events (ADEs) than did chart review but many more ADEs than did stimulated voluntary report. The overlap among the ADEs identified using different methods was small, suggesting that the incidence of ADEs may be higher than previously reported and that different detection methods capture different events. The computer based monitoring system represents an

> efficient approach for measuring ADE frequency and gauging the effectiveness of ADE prevention programs." [400]

The previous paragraphs have focused on the technical problems associated with obtaining accurate assessments of the participation ratio; the total number of contributed reports divided by the total expected frequency of incidents. However, it is important not to underestimate the managerial consequences of such work. The process of obtaining a more accurate assessment of underlying incidents can itself trigger enormous changes within an organisation:

> "In February 1999 a urologist at the Sturdy Memorial Hospital in Attleboro, Massachusetts, requested a retrospective review of a 1996 biopsy result because of the patient's clinical course and the results of a biopsy in 1999. The review revealed that the 1996 report was incorrect. The urologist and pathologist (neither of whom was responsible for the 1996 reading) implemented appropriate management for the affected patient.
>
> When they discovered a second misread prostate biopsy from the same period the urologist and pathologist became concerned that the frequency of these errors was higher than "expected". Fears about malpractice suits and damaged reputations emerged... Ultimately, the medical director thought that all the prostate biopsies performed during 1995-7, the period of tenure of the clinicians associated with the two errors, should be reviewed... During the review we wondered about any requirements to report to regulatory agencies. Our lawyers told us we had no obligation to report this kind of error... We decided to report our initial findings to the Department of Public Health and the Board of Registration in Medicine. In total 20 of the 279 prostate biopsies from 1995-7 were in error. The urologists caring for these 20 patients were told of the changes in the biopsy interpretations, and it was agreed that the urologists would contact each patient and recommend appropriate evaluation and treatment. Although they agreed with this plan, the urologists were worried about potential lawsuits, damage to their reputations, and the stress of difficult meetings with the patients and their families.
>
> When the process of notifying the patients started, the hospital president realised that questions about the validity of other biopsies would be raised even though there was no clinical evidence to raise such concern. She thought that all should be reviewed... About 6000 biopsies would have to be reread, and we needed help. Inquiries to the professional pathology bodies were disappointing: not only did we receive little assistance, but we were routinely asked why we wanted to expose more errors..." [682]

Many of the ethical worries that affected the physicians in this case, stemmed from the voluntary nature of incident reporting within their profession, Mandatory reporting systems offer alternative means of addressing the problems of under-reporting. They simplify the previous dilemma at the cost of restricting an individual's freedom to choose whether or not to report a particular incident.

5.1.2 Mandatory Reporting

The UK Air Accident Investigation Branch has published formal accident reports to disseminate the lessons that have been learned from air proximity warnings. Individuals are obliged to report these near-miss incidents in the same manner that they are obliged to report accidents. This obligation to report is enshrined within the 'Duty to furnish information relating to accidents and incidents' paragraphs of the Civil Aviation (Investigation of Air Accidents and Incidents) Regulations 1996:

> "5.(1) Where an accident or a serious incident occurs in respect of which... the Chief Inspector is required to carry out, or to cause an Inspector to carry out, an investigation, the relevant person and, in the case of an accident or a serious incident occurring on or adjacent to an aerodrome, the aerodrome authority shall forthwith give notice thereof to the Chief Inspector by the quickest means of communication available and, in the case of an accident occurring in or over the United Kingdom, shall also notify forthwith a police officer for the area where the accident occurred of the accident and of the place where it occurred." [11]

These regulations, in turn, depend upon definitions of accidents and incidents. Section 1.2.2 reviewed a number of different techniques that have been used to distinguish between these different classes of occurrence. However, the UK Civil Aviation Regulations follow the approach proposed in ICAO Annex 13:

> 'accident' means an occurrence associated with the operation of an aircraft which takes place between the time any person boards the aircraft with the intention of flight until such time as all such persons have disembarked, in which
>
> - (a) a person suffers a fatal or serious injury as a result of- -being in or upon the aircraft -direct contact with any part of the aircraft, including parts which have been detached from the aircraft, or -direct exposure to jet blast, except when the injuries are from natural causes, self-inflicted or inflicted by other persons, or when the injuries are to stowaways hiding outside the areas normally available to the passengers and crew, or
> - (b) the aircraft sustains damage or structural failure which adversely affects the structural strength, performance or flight characteristics of the aircraft, and would normally require major repair or replacement of the affected component, except for engine failure or damage, when the damage is limited to the engine, its cowlings or accessories; or for damage limited to propellers, wing tips, antennas, tyres, brakes, fairings, small dents or puncture holes in the aircraft skin; or
> - (c) the aircraft is missing or is completely inaccessible...
>
> ... 'serious incident' means an incident involving circumstances indicating that an accident nearly occurred. [11]

These regulations illustrate the way in which legal obligations can be placed upon operators so that they are *required* to report certain categories of near-miss incidents. There are examples of similar mandatory systems in other domains. For example, the recent UK National Health Service report entitled 'An Organisation with a Memory' proposed a national mandatory reporting scheme for adverse health events, and specified near misses, based on standardised local reporting systems [633]. There are, however, mixed views about the effectiveness of such systems. For example, the Safe Medical Devices Act of 1990 requires that healthcare facilities and manufacturers must report serious injury or illness related to the failure or misuse of specific medical devices. However, Cohen has argued that:

> ...this federal act has been unsuccessful in gaining compliance with reporting requirements for user error. Furthermore, little action is taken unless significant numbers of harmful errors have been reported. Some states also have mandatory reporting programmes for error resulting in serious patient harm. Yet this information is used almost exclusively to punish individual practitioners or healthcare organisations. There is little analysis of the systems causes of error, and the information is rarely used to warn others about the potential for similar errors. ...non-punitive and confidential voluntary reporting programmes provide more useful information about errors and their causes than mandatory reporting programmes. A major reason is that voluntary programmes provide frontline practitioners with the opportunity to tell the complete story without fear of retribution..." [171]

Many of Cohen's criticisms seem to focus on ways in which mandatory systems have been used, or 'abused', by those who operate them. Very few of his adverse comments directly stem from weaknesses in mandatory systems. There are, however, strong concerns about the enforcement of mandatory systems. Clearly, if an individual or group have suppressed information about an incident then others within the organisation must be in a position to detect it if any form of action is to be taken. If an individual fails to report a mandatory occurrence then they run the risk that one of their colleagues may also detect and submit information about an incident. Follow-up investigations might then centre on the reasons why the first operator failed to provide any notification of the adverse event. Alternatively, incidents can come to light through the post hoc review of logs and records. This approach relies upon techniques that are similar to the exhaustive analysis that has been used to identify background incident rates, and thereby derive reporting quotas. Irrespective of the source of such information, there remains the problem of determining what disciplinary action should be taken when individuals fail to report mandatory incidents. Typically, this depends upon the severity of the incident being considered and upon whether the individual had a clear appreciation of that severity. For instance, if the incident occurred during a period of high workload, it may not be certain that the operator did actually detect the adverse event. Even if they did detect it, the pressure of other duties may have prevented them from reporting it. High workload may even contribute to individuals forgetting about low-criticality occurrences [863].

It is also important to stress that mandatory reporting systems need not be based upon the regulatory or legislative model. For example, they can be integrated into everyday working practices. Individuals and groups may be required to fill in an occurrence reporting form after every procedure, operation or shift. In most instances, the form records the fact that no incident had occurred. However, the insistence that such a form is routinely completed may help to raise the prominence of the system. This approach can encourage greater participation in the reporting system. The drawbacks are also readily apparent. There is no guarantee that the routine completion of an incident reporting form will have a positive impact upon reporting behaviour. There is also the danger that the additional workload may alienate staff from using the system when incidents do occur. Some of these objectives can be addressed by integrating the routine reporting activity into other everyday tasks. For example, the completion of a medical incident reporting form could be intergrated with minimal overhead into existing patient documentation. Barach and Small provide a more optimistic assessment of the utility of both mandatory and voluntary systems:

> "Mature safety cultures are driven by forces external and internal to industries, and over time these forces nourish voluntarism and reporting of near misses. Furthermore, rapidly improving technology and information systems enable wider monitoring and public awareness of adverse outcomes in open systems. These developments diminish distinctions between mandatory and voluntary behaviour." [66]

The previous paragraphs document the expressed intuitions of practitioners who are developing incident reporting systems within their particular domains. As with many other aspects of incident reporting, there is a pressing need for more reliable data to back-up these assertions about the impact that different voluntary and mandatory approaches will have upon the notification of information about adverse occurrences.

5.1.3 Special Initiatives

Previous sections have argued that voluntary reporting systems suffer from considerable problems of under-reporting. Mandatory systems can address some of these problems, however, they can alienate some members of staff and have not been universally successful. Special initiatives provide an alternate incident reporting technique that can be used to address under-reporting. At their simplest, these initiatives may simply be implemented through simple questionnaires that directly poll staff about incidents and issues that have occurred to them in recent months. This approach has the benefit that all staff may be called on to participate at the same time and in a confidential manner through the return of a simple form. For instance, Sexton, Thomas and Helmreich [735] have exploited this approach to examine more general attitudes to error within the medical profession. Their study prompted returns from 851 operating room staff and 182 intensive care workers. The data that can be obtained using such questionnaires is very different from the more focused information that is provided by conventional incident reporting systems. However, it would have taken many years to elicit the same number of response through more conventional incident reporting systems. More importantly such initiatives can be used to examine the reasons why particular groups *fail* to participate in incident reporting systems even though they may acknowledge that these systems form a valuable part of any safety system:

> "Over 94% of intensive care staff disagreed with the statement 'Errors committed during patient management are not important, as long as the patient improves'. A further 90% believed that 'a confidential reporting system that documents medical errors is important for patient safety'. Over 80% of intensive care staff reported that the culture in their unit makes it easy to ask questions when there is something they don't understand (this is undoubtedly related to the high endorsement of flat management hierarchies in the unit). One out of three intensive care respondents did not acknowledge that they make errors. Over half report that decision making should include more team member input.
>
> More than half of the respondents reported that they find it difficult to discuss mistakes, and several barriers to discussing error were acknowledged. The 182 staff in intensive care reported that many errors are neither acknowledged nor discussed because of personal reputation (76%), the threat of malpractice suits (71%), high expectations of the patients' family or society (68%), and possible disciplinary actions by licensing boards (64%), threat to job security (63%), and expectations or egos of other team members (61% and 60%). The most common recommendation for improving patient safety in the intensive care unit was to acquire more staff to handle the present workload, whereas the most common recommendation in the operating theatre was to improve communication." [735]

As mentioned, questionnaire based techniques are qualitatively different from other forms of incident reporting. They help to reveal general attitudes rather than specific information about particular adverse occurrences. On the other hand, such initiatives are deeply revealing about the attitudes to error that chapter 3 has argued to be significant causes of more 'systemic' failures.

Questionnaire-based techniques can also be used to examine the biases that can skew the under-reporting of particular sorts of incidents. For example, there is a greater danger that low-consequence incidents, well-known problems will not be routinely reported. Martin, Kapoor, Wilton and Mann provide valuable insights into the nature of these problems in the medical domain:

> Data on side effects of newly launched drugs are limited,1 highlighting the need for effective post-marketing surveillance. An inverted black triangle on product literature identifies new products. Suspected adverse reactions to these drugs, however minor, should be reported to the Committee on Safety of Medicines through the yellow card scheme. Adverse reactions are Adverse reactions are underreported, and few doctors in the United Kingdom know the meaning of the 'black triangle' symbol. We assessed the degree of underreporting of suspected adverse reactions to new drugs in general practice and determined if reporting varied when reactions were severe or previously unrecognised.
>
> There were 3045 events (in 2034 patients) reported as suspected adverse reactions on the green forms during the 10 studies. General practitioners indicated that they had reported 275 (9.0%; 95% confidence interval 8.0% to 10.0%) of these reactions to the Committee on Safety of Medicines: reporting was highest for serious unlabelled reactions (26/81; 32.1 %) and lowest for non-serious labelled reactions (94/1443; 6.5 %). Serious unlabelled and non-serious unlabelled reactions were significantly more likely to be reported than were non-serious labelled reactions. According to general practitioners' responses, the proportion of serious labelled reactions also reported on yellow cards (7/64; 10.9%) was only slightly greater than that of non-serious labelled reactions." [522]

The strength of this work is that Martin et al show how it is possible, in certain circumstances, to obtain objective data about the extent and nature of the under-reporting problem. The 'green forms' mentioned in the previous quotation were questionnaires that had been distributed by the researchers to general practitioners. These voluntary returns were then correlated against self-reported mandatory returns to the Committee on Safety of Medicines. Of course, even these results are subject to recall and reporting biases but they do indicate how focused initiatives can be used to elicit more information about the nature and extent of under-reporting [504].

Questionnaires are not the only form of special initiative that can supplement more conventional or general forms of incident reporting. In particular, issue based reporting systems have been used to overcome the under-reporting of particular critical occurrences. For example, many organisations established incident reporting systems that were specifically intended to provide information relating to potential problems during the millennium period. The UK MDA implemented an Adverse Incident Tracking System. This was intended to provide the NHS with information on issues involving medical devices during the period 23rd December 1999 to 10th January 2000. This database supplemented the MDA's normal Hazard and Safety Notice systems.

Issue based incident reporting systems can also be used to make inferences about the background rate of contributions. The difference between the reporting frequency before the initiative and after the initiative can be used both to gauge the success of any focus on particular issues and to provide a more general measure of under-reporting. For example, the UK Meteorological Office operated the European Turbulent Wake Monitoring Scheme between 1995-1999. This was set up by the European Commission because the separation minima at airports take no account of existing meteorological conditions. They are simply calculated using the weight of the aircraft. Under favourable meteorological conditions, however, it may be possible to reduce these minima if a wake vortex is less likely to occur. The intention was to create and maintain a database of wake vortex incident reports with associated meteorological data. Researchers and aviation companies could then use this data to better understand wake vortex behaviour. A voluntary incident reporting system was chosen as a means of compiling this data because fully equipped meteorological monitoring systems cost up to $1m for a single airport. The initiative was intended to address under-reporting problems because the UK was the only European country to regularly monitor wake vortex incidents. However, over 90% of reported encounters in this existing system took place around Heathrow airport. There was "clearly a need for data from airports with a diverse range of runway configurations, meteorological phenomena and capacity in order to assess the global problem" [547]. It can, however, be argued that this scheme illustrates some of the limitations of such focused initiatives. The system ceased to record further data once the initial funding from the European Commission had run out.

5.2 Encouraging the Detection of Incidents

Previous sections have argued that under-reporting continues to be a significant problem for many incident reporting systems. Mandatory participation provides a potential solution but also raises further pragmatic and ethical problems. Special reporting initiatives can be used to assess the scope and nature of the under-reporting problem. However, pro-active questionnaires and systems that are focused on specific types of incidents suffer from different forms of reporting bias. It can also be difficult to sustain high levels of participation in special reporting initiatives. The following pages reviews manual and automated techniques that can be used to combating the problem of under-reporting. These techniques are intended to support the more general class of voluntary incident reporting systems introduced in Chapter 4, rather than special purpose or mandatory systems.

5.2.1 Automated Detection

This section focuses on automated techniques that reduce the need for individuals and groups to explicitly contribute occurrence reports. As we shall see, however, there are a number of technical and organisational concerns that can complicate the introduction and application of these systems. These include the alienation and lack of trust that can emerge when automated systems either fail to detect incidents or, conversely, when systems erroneously spot incidents that did not threaten safety. There are also concerns that the introduction of such systems represents an unwarranted intrusion into the working lives of those whose actions are being monitored.

Trust and Acceptance

This book has primarily focused on incidents that are detected by human operators. As reporting systems become more established, however, it is also possible to use automated tools to supplement this source. However, different industries offer different opportunities for the automated detection of critical incidents. Previous sections have described how simple database tools can be used to search through electronic patient records to support manual chart monitoring techniques. Air Traffic Service networks provide ground and airborne systems such as ground proximity warning systems (GPWS), minimum safe altitude warning (MSAW) systems, short-term conflict alerts (STCA), aircraft proximity warning (APW) and aircraft collision avoidance system (ACAS).

It is possible to identify two different roles for the systems that support the automated detection of adverse occurrences:

- *on-line alerts.* Automated systems can warn operators about a safety occurrence that is taking place or about the potential for a more severe occurrence. They can be used to monitor and trigger occurrence-reporting procedures when they automatically detect that certain adverse circumstances have occurred. For example, workers or their supervisors might be expected to make a preliminary report whenever a warning is generated. As we shall see, problems arise when these on-line systems incorrectly diagnose that an incident has occurred. There is a paradoxical danger that such alarms may trigger genuine events as operators struggle to dismiss unwanted warnings;

- *post hoc monitoring.* Automated systems can also be used off-line to search for adverse occurrences. This approach is more suitable when the outcome of an event may not be known for some time after an initial procedure has been completed. For instance, medical incident reporting systems may have to assess the success or failure of an intervention in terms of the patient's quality of life months or even years after they have been discharged. Although there are a number of potential problems in mixing safety issues with more general process improvement concerns, there is an increasing move towards this type of incident reporting architecture [453].

As mentioned, the degree of sophistication in the automation that is available to detect potential incidents varies widely from industry to industry. The development of this technology depends both upon the complexity of the application that is being controlled. For example the ability to monitor pilot actions might be interpreted as a by-product of the development more advanced control systems. The development of automated detection technology also depends upon the consequences of failure and the severity of the perceived threat. Although not directly a safety concern, this can be illustrated by recent initiatives to improve the monitoring of security incidents involving US Department of Defence Computers. These represent instances of the malicious failures described in Chapter 2:

> Rapid detection and reaction capabilities are essential to effective incident response. Defence is installing devices at numerous military sites to automatically monitor attacks on its computer systems. For example, the Air Force has a project underway called Automated Security Incident Measurement (ASIM) which is designed to measure the level of unauthorised activity against its systems. Under this project, several automated tools are used to examine network activity and detect and identify unusual network events, for example, Internet addresses not normally expected to access Defence computers. These tools have been installed at only 36 of the 108 Air Force installations around the world. Selection of these installations was based on the sensitivity of the information, known system vulnerabilities, and past hacker activity. ASIM is analysed by personnel responsible for securing the installation's network. Data is also centrally analysed at the Air Force Information Warfare Center (AFIWC) in San Antonio, Texas. Air Force officials at AFIWC and at Rome Laboratory told us that ASIM has been extremely useful in detecting attacks on Air Force systems. They added, however, that as currently configured, ASIM information is only accumulated and automatically analysed nightly. As a result, a delay occurs between the time an incident occurs and the time when ASIM provides information on the incident. They also stated that ASIM is currently configured for selected operating systems and, therefore, cannot detect activity on all Air Force computer systems... DISA officials told us that although the services' automated detection devices are good tools, they need to be refined to allow Defence to detect unauthorised activity as it is occurring. DISA's Defensive Information Warfare Management Plan provides

> information on new or improved technology and programs planned for the next 1 to 5 years." [761]

This quotation describes military systems that are intended to automatically detect external threats to computer system security. Entirely different issues are raised when automated systems are employed to detect human 'error' and system 'failure' that stems from non-malicious acts within an organisation. In particular, the effective use of automated monitoring devices is not simply determined by technology sophistication. It is also profoundly determined by social and managerial issues. Irrespective of the technology that is being used, it is critical that automated monitoring tools gain staff acceptance.

Trust and Acceptance

The importance of staff acceptance of automated monitoring devices cannot be underemphasised. The action of trades unions and other forms of worker representation can block the introduction and use of this technology for many years. Driver monitoring systems on UK railways provide a good illustration of this point. In 1999, Her Majesty's Railway Inspectorate (HMRI) issued a report that analysed the management systems which were intended to record and assess incidents in which drivers had passed signals 'at danger'. The number of these incidents that were reported in the UK gradually fell during the 1990's. However, it has levelled out in recent years: 944 in 1991/92, 593 in 1997/98, 643 in 1998/99 These incidents have also led to a number of high-profile accidents. The collision at Watford South Junction on 8 August 1996 caused the death of one passenger. The accident at Southall on 19 September 1997 was also caused by a signal being passed at danger. As a result of these accidents, plans were developed for the introduction of the Train Protection and Warning System (TPWS). This is intended to mitigate the effect of such incidents by warning the driver and ultimately by braking the train at junctions, single lines and 'unusual' train movements, see Chapter 3. However, the cost and complexity of such equipment has delayed its introduction. As an interim measure a range of Driver's Reminder Appliances (DRA) have now been fitted to most driving cabs. These have the limited role of reminding a driver of the current signal when they are stopped at a station with the starting signal at danger. Without more advanced protection systems, an argument was made for more closely monitoring driver behaviour. This was based on the idea that human factors problems could be addressed through remedial training and supervision if it was possible to identify those drivers who were most likely to pass signals at danger. The HMRI report reviewed piecemeal progress towards the introduction of driver monitoring equipment that was intended to make this possible:

> "Recommendation 9 of the HMRI report into the accident at Watford South Junction was to North London Railways (now Silverlink) to extend the use of on train data recorders to monitor driving technique. Although the number of trains fitted with the equipment is still less than 20% of the total, the number is increasing rapidly so other Train Operating Companys should be making use of it for unobtrusively monitoring driver performance. Thameslink was not doing so (although was to start) and neither was Connex South Central, although it is acknowledged that most of their fleet is not fitted with the equipment." [349]

More recently, the action plan to implement the recommendations of the Southall accident report again included steps to extend the CIRAS voluntary incident reporting system *and* automated monitoring equipment. The explicit reference to the drivers union ASLEF (Associated Society of Locomotive Engineers and Firemen) is instructive:

> "Evidence (of driver involvement in 'signal passed at danger' incidents) should include that to be provided by CIRAS and from On-Train Data Recorders used to monitor driver behaviour. ASLEF in particular should give their full support to such an initiative." [317]

This comment about the need for ASLEF support is important because it reveals the HMRI's sensitivity to workers' concerns about the introduction of these automated sensing systems into the cabs. It is important to emphasise that these concerns again rest on a justified fear of retribution that affects all forms of incident reporting. These fears are exacerbated by a number of additional factors. As we shall see, the sensitivity of these devices can lead to false readings that might, in turn, trigger unwarranted accusations of poor performance and error. The piecemeal and delayed introduction of these systems may mean lead to inequitable treatment. Some driver errors may be 'caught' by these systems whilst others may go undetected because no equipment is installed. Automated equipment triggers the detection of some errors, however, it often fails to capture the 'mitigating' factors that can excuse certain violations. Finally, such automated systems address the observed consequences of deeper systemic failures, including poor signal placement, that actually *cause* the failure in the first place [732] Many of these concerns do not stem from the ethical involved in introducing automated monitoring equipment. Instead, they centre on the ways in management will use the data that is collected by these systems. Such concerns were touched upon by an earlier enquiry. The following quotation is revealing because it probes the limits of a 'no blame' culture. The ambivalent position of the regulators again illustrates how pragmatics lead to what we have termed a 'proportionate blame' approach. The report does not criticise the use of SPAD reports, from either manual or automated sources, within a company's disciplinary procedures.

> "All Train Operating Companys (TOCs) visited have specially monitored driver procedures in place to assign drivers to categories dependent on their incident history. However, the application of these procedures (mandated by Group Standard GO/RT3251) varies widely between TOCs. For example Connex South Central at some drivers' depots allocated all drivers to one of the 'at risk' categories. West Anglia Great Northern Railway's (WAGN) procedure appeared to give rise to too much scope for management discretion in reducing the status of a driver from 'incident prone' to 'normal'. Since, HMRI's inspection, WAGN are revising their procedure.
>
> Drivers in higher risk categories are intended to be subject to a greater number of assessment rides focusing on identified weaknesses, but it is questionable whether these are always achieved in practice. Generally, these are managed by individual drivers' depots, but it would be more satisfactory for this to be monitored centrally within TOCs to ensure that the extra assessments are actually carried out and that they address any identified weaknesses in competence.
>
> There must be adequate procedures for removing a driver from driving duties in the event of their SPAD record not improving despite further training and assessment. Some TOCs use the disciplinary procedure, but the key requirement for TOC managements following incidents is to ensure any deficiencies in competence are identified and robustly addressed by means of further training, if necessary, and competence assessment..." [349]

This quotation also illustrates how inconsistent management practices can lead to different companies reacting in different ways to drivers committing the same 'errors' on the same piece of track. For the proponents of no-blame cultures, it is salutary to note that the HMRI found improved safety records in those companies that adopted a 'hard-line' approach to SPADS. Without automated equipment and lacking any details of the procedures used to elicit information about SPAD incidents within those companies, it remains likely that the 'hard-line' approach simply dissuade drivers from contributing information about these adverse occurrences:

> "The version of GO/RT3252 in use at the time of HMRI's inspection required that when a driver had had three SPAD incidents, they were only to continue on driving duties if there was a written justification for doing so. This was not always found to be the case. Some TOCs were found to take a relatively hard line and removed any driver automatically from driving duties at the third SPAD incident, whereas others did not. It could be significant that those TOCs which were found to take a hard line in this area appeared to have better SPAD records than others, and this may lead drivers to adopt the required defensive driving approach. The new version of GO/RT3252, revised since HMRI's inspection, focuses more on the identification and rectification of competence weaknesses which lead to SPAD incidents. " [349]

This section has argued that there are a number of reasons why workers may distrust automated incident detection systems. These include concerns about the way in which information from these systems will be used during any subsequent disciplinary hearings. Trust in automated detection systems can also be eroded for technical reasons. These include the problems that reduce the signal to noise ratio associated with particular warnings. In particular, there can be problems with missed incidents and false alarms.

Missed Incidents and False Alarms

Chapter 3 argued that environmental features can prevent operators from accurately perceiving important properties of their environment. Table 5.2.1 was used to show how a signal may or may not be present. If the signal is present then either the operator may detect it, in which case they have achieved a 'hit', or they may fail to detect the signal, this results in a 'miss' in Table 5.2.1. If the signal is absent and the user detects it then this results in a false alarm. However, if they do not detect a signal then this represents a correct rejection. These same distinctions apply both to the human detection of signals or warnings in their environment and to the automated detection of critical incidents. For instance, if an automated system detects a signal, that is to say an incident, when none is present then this will generate a false alarm. Conversely, if an incident did occur and was detected then this represents a 'hit' by the detection equipment. A 'miss' occurs if an incident took place but was not detected. A correct rejection takes place when the system successfully finds that no incident has occurred. Wiener summarises the technical problems that emerge from this analysis:

Response		State of the World	
		Signal	Noise
	Yes	Hit	False Alarm
	No	Miss	Correct Rejection

Table 5.1: Outcomes from Signal Detection

> "In any warning system, one can expect false alarms and missed critical signals, and the designer must design the filter logic to strike a balance. If the system is deigned to be 'sensitive', that is to have a high detection rate, then it will hive a high false alarm rate, and vice versa. There is no perfect system that can detect all true events and filter out all false events." [864]

The problems that this creates are illustrated by the strengths and weaknesses of Traffic Alert and Collision Avoidance System (TCAS) II. In 1987, the FAA mandated the installation of TCSII equipment on all airliners by the end of 1993. In general terms, this equipment provides two levels of warning. The first is issued 45 seconds before the predicted point of closest approach This consists of a display that present the distance and bearing of the target aircraft. Between 20 and 25 seconds before the predicted point of closest approach, a resolution advisory is sounded. This, typically, requests a vertical manouvre to increase separation. It is clear that TCAS II has saved many lives, however, initial implementations raised numerous problems. In particular, the sensitivity that Wiener argues is essential to detect potential incidents can also add to crew workload when a situation is already being resolved:

> "...we received two TCAS II-advisories, corresponding to departures. The departures are cleared to 10,000 feet, [and] arrivals...[at] 11,000 feet. The TCAS II reacted to the closure rate of the departing aircraft and our inbound flight. [The] RA was ignored as traffic was in sight. The real problem is that the TCAS II alert caused such a distraction in the cockpit that two or more radio calls from Approach Control were missed." [546]

The conditions that lead to spurious alarms are hard to anticipate. For example, some relate to technical failures in the manner in which aircraft altitude data is acquired from the Mode C function of the aircraft's radar transponder. Should Mode C even temporarily provide erroneous altitude information, an erroneous TCAS II Resolution Advisory command to climb or descend may result [546]. Other false alarms can be generated by local features. For instance, Billings cites numerous spurious warnings at particular airports including Orange County California and Dallas Fort-Worth Texas. He argues that such missed incidents and false alarms have a considerable impact upon the behaviour of system operators. Early versions of the Ground Proximity Warning System (GPWS) were so unreliable that crews ignored or disabled them. Such actions indirectly led to accidents at Kaysville Utah (1977) and Pensacola, Florida (1978) [82]. One large commercial airline discovered 247 (73%) spurious alarms amongst a total of 339 GPWS alerts in a twelve month period.

In passing it is worth mentioning that incident detection systems, such as TCAS, can influence operator behaviour in ways that threaten rather than preserve the safety of an application. For example, pilots may often perform violent descents or ascents in response to an advisory. These manouvres may, in turn, raise TCAS advisories on other aircraft. The knock-on effects of this behaviour is to significantly increase the burdens on Air Traffic Control officers. This creates a paradoxical situation in which the introduction of incident monitoring systems may actually contribute to an increase in the adverse occurrences that they were intended to detect:

> "Air carrier (X) was inbound on the...STAR level at 10,000 feet. Under my control, air carrier (Y) departed...on the...SID, climbing to [an] assigned altitude of 9,000 feet. Approximately 14 miles SW...I issued traffic to air carrier (X) that air carrier (Y) was leveling at 9,000. Air carrier (X) responded after a few seconds that they were descending. I again told air carrier (X) to maintain 10,000 feet. Air carrier (X) responded 'OK, we've got an alert saying go down'. Simultaneously, air carrier (Y) was getting an alert to climb. They both followed the TCAS II advisorys and almost collided. Later, [the pilot of air carrier (X)] ...indicated [that] his TCAS II was showing zero separation. They passed in the clouds without seeing each other. When pilots start taking evasive action, our equipment cannot update quickly enough for the controller to help. Both aircraft were issued traffic as prescribed by our handbook (merging target procedures). [Air Carrier] Company directives, I'm told, dictate that pilots must respond/follow the TCAS II alert advisories." (ACN 224796) [546]

Currently, TCAS II advisorys do not automatically trigger the generation of an incident report. This is best explained in terms of a further paradox. In order for monitoring systems to provide real-time warnings to operating personnel, they must be so sensitive to potential incidents that they may generate a number of spurious warnings. This high number of spurious warnings imposes too high a workload for each alarm to be individually investigated and reported. As a result, the warnings provided by such systems are often filtered by informal operating practices so that only a small proportion of the detected events are notified to a reporting system. For example, the initial installation of TCAS II led to a high level of ASRS reports. There were 1,996 TCAS related submissions between 1988 and 1992 alone.

Limited Views of Causation

A number of safety concerns emerge from the integration of automated incident detection systems into complex working environments. The previous section argued that this can, itself, jeopardise safety if spurious alarms cause deviations from normal operating practise or if individuals respond in unpredictable ways. There are further concerns that relate more narrowly to the practice of incident reporting. In particular, there is a danger that operators will come to rely on incident detection systems. For example, the 'security' provided by TCAS can indirectly degrade other forms of vigilance:

> "I was training a developmental [controller] on Arrival Control. We had an air taxi (X) for sequence to visual approach Runway 15. The developmental pointed out aircraft (Y) [to air taxi (X)] and the pilot responded, 'Is he following someone out there at 800 feet?' The developmental was going to clear him for the visual approach when I stopped him and asked [the pilot of air taxi (X)]...if he had aircraft (Y) in sight. He said not visually,

> but had him on TCAS II. This seems to be happening more and more...It appears [that pilots]...are using TCAS II instead of looking out the window. As an air traffic controller I cannot have pilots using TCAS for visual separation to maintain spacing (as on one occurrence a crew offered to do). There is no TCAS II separation." ([546], ACN 202301)

There are a number of reasons for this concern. In particular, systems such as TCAS are intended to alert crews to adverse circumstances that should not occur during normal operation. They form part of a safety net that is intended to save personnel from the adverse consequences of those situations. If they are assimilated into everyday operation practices then the additional assurance provided by those systems will be lost.

The particular consequences for incident reporting are that the (ab)use of automated detection systems makes it less likely that personnel will explore the underlying causes of the alarm that they have experienced. This is important because technologies, such as TCASII, minimum safe altitude warning (MSAW) systems, short-term conflict alert (STCA), can be used to trigger investigations that stand some chance of uncovering the deeper systemic issues that exposed users' to danger in the first place. Both Perrow [675] and Reason [701] warn if these underlying causes are not addressed then it is likely that our defences will fail at some point in the future. Each TCAS warning in aviation or SPAD in the rail industry not only warns the individual pilot or driver, it should also be a warning to the industry as a whole.

A number of pragmatic issues limit the amount of information that can be obtained from automated incident detection systems. For instance, TCAS II provides limited data about aircraft separation. It does not provide a 'complete' account of the causal factors and influences that led to the loss of separation. Automated recording equipment can provide more detailed insights into the course of an incident. For example, digital flight data recorders provide information about a failure to fly a stabilised approach, about engine overspeed and about an excessive rate of descent. Over time such data can be collated to provide an overview of common problems, for example repeated overspeeds by several pilots when landing at a particular airport [342]. However, it is important to acknowledge the limitations of the data that can currently be captured. For instance, it is not possible to use digital flight data to determine what caused the specific incidents that are recorded. It may be due to pilot 'error', to air traffic control restrictions, to adverse meteorological conditions etc. In other words, the information that is elicited by automated systems currently only acts as a trigger for further investigation. In this respect, it is no different from the trigger that is provided by the manual detection and contribution of incident reports.

The need to supplement the information that is obtained by automated resources has focused into a debate about whether or not video recorders should routinely be used to supplement the cockpit voice recorders on aircraft. Jim Hall, Chairman of the National Transportation Safety Board (NTSB) gave the following testimony before the Subcommittee on Aviation Committee on Transportation and Infrastructure in the House of Representatives:

> "...the Safety Board's investigation into several recent crashes has highlighted the need for recording images of the cockpit environment. The Safety Board believes that the availability of electronic cockpit imagery would help resolve issues surrounding flight crew actions in the cockpit. For example, it would tell us which pilot was at the controls, what controls were being manipulated, pilot inputs to instruments (i.e., switches or circuit breakers), or what information was on the video displays (i.e., the display screens and weather radar). Video recorders would also provide crucial information about the circumstances and physical conditions in the cockpit that are simply not available to investigators, despite the availability of modern cockpit voice recorders (CVRs) and 100-parameter digital flight data recorders (DFDRs).
>
> The Safety Board first discussed the need for video recording the cockpit environment in its report of the September 1989 incident involving USAir flight 105, a Boeing 737, at Kansas City, Missouri. In that report, we recognised that while desirable, it was not yet feasible... Electronic recording of images in the cockpit is now both technologically and economically viable, and solid state memory devices can now capture vast amounts of audio, video and other electronic data. ...the Safety Board is extremely sensitive to

> the privacy concerns that the pilot associations and others have expressed with respect to recording images of flight crews. As you know, the Board's reauthorisation passed by this Chamber would require that the same protections already in place for CVRs be extended to image recorders in all modes of transportation. Under those provisions, a cockpit image recording could never be publicly released. Even where monitoring has been allowed there is resentment towards certain technologies." [302]

This quotation does acknowledge the concerns that commercial airline pilots feel about the introduction of such systems. These concerns were intensified when several media organisations broadcast the final minutes of the cockpit voice recorder during the Cali crash. Although this would not have been allowed under US or Canadian legislation, there was no provision to prevent the release of such material in Columbia at the time of the accident. There is also the perception amongst pilots that such video equipment is being introduced to satisfy public perceptions about the utility of recordings and that these perceptions may not, in fact, be justified. This argument has considerable strength. Chapter 3 noted the difficulty of interpreting intention and cause from video recordings of users who commit 'everyday', non-safety critical errors. These difficulties would be considerably greater in the aftermath of an accident.

Fortunately, near-miss incidents offer alternative means of eliciting additional information to support the output of automated monitoring equipment. Several major airlines have now installed Air Data Acquisition Systems (ADAS) that record a range of information that is typically already recorded by the digital flight data recorders ('black box' recorders). For example, British Airways currently supplement their air safety reporting programme with data collected from their SESMA flight data recorders [658, 659]. ADAS information can be routinely monitored to detect whether certain triggering conditions occur during otherwise normal operation. These triggering conditions include warnings from GPWS, TCAS, stall protection systems etc. They can include attitude transgressions, such as overbanks, or the transgression of speed limits, such as flap overspeeds. They may also include incidents involving extreme g-loads or prolonged flares. If a trigger occurs then the airlines' flight data analysts may interview the crew. Klampfer and Grote used this technique to analyse 71 incidents within a commercial fleet [445]. 48 of the incidents involved A320 aircraft, 18 involved the MD11 and the rest were from a variety of other aircraft. This data revealed that 29% of incidents involved speed violations, analysts included underspeed and overspeed conditions in this category. 19% of incidents involved unstable approaches. 11% involved prolonged flares. 10% involved low go arounds. 10% of all incidents were triggered by the automated monitoring systems mentioned in previous sections. 10% of the incidents involved attitude violations. 8% involved excess g-loads. 3% of the events could not be classified according to these general categories/ The interview data was examined together with the triggering information from the ADAS system. A causal analysis identified that direct human errors contributed to 40% of all incidents. These errors included poor situation awareness and a lack of crew coordination. Human influences contributed to 31% of the incidents. These are classified as actions that are not, of themselves, incorrect but which contributed to or exacerbated the consequences of an incident. This is perhaps the most difficult of Klampfer and Grote's categories; it includes mental overload and routine action as contributory causes. Their analysis also identified that 16% of incidents were caused by environmental factors, including air traffic control 'failure'. Only 11% were caused by technical failures, including poor meteorological information. The remaining 2% were unclassified.

Although it is possible to question the taxonomy that Klampfer and Grote use in their analysis, it is important to recognise the benefits that their pioneering use of autoated detection and manual investigation can provide. It can be used in a non-punitive manner to examine common causes between a number of incidents. It also provides important checks and balances to the work of the incident investigator who might otherwise form a number of unwarranted assumptions on the basis of limited ADAs data. There are also a number of unexpected benefits. In particular, this technique can be used to probe for a potential, unreported loss of situation awareness or long-term consequences of adverse occurrences when aircrew recollections differ significantly from the information recorded by the ADAS infrastructure.

Chapters 10 and Chapter 15 will look at conventional tools, including relational databases, and more advanced techniques, such as case based reasoning, that support the off-line monitoring of incident reports. In contrast, this section has focused on systems that support the on-line, or run-time, detection of adverse occurrences. These systems offer a number of important advantages. In particular, they can help operators to avoid a potential incident or mitigate the consequences once an incident has taken place. The same systems can also be used to trigger further causal investigations after an event has occurred. A particular concern is that most regulatory and commercial organisations focus on the former role of automated detection systems. They often miss the opportunity to address the causes of those incidents that are reported by automated detection equipment. As a result, latent weaknesses become embedded in systems that rely upon detection equipment as primary rather than a secondary defence mechanism.

5.2.2 Manual Detection

The previous section has identified ways in which automated systems can be used to monitor operating logs in order to detect potential incidents and accidents. In contrast, the following paragraphs focus on techniques that are intended to encourage individuals and groups of workers to manually submit safety-related information. A number of general guidelines are followed by a more sustained and detailed analysis of the different forms that can be used to elicit critical data.

A Reporting Culture

Previous sections have argued that a proportionate blame approach is an important prerequisite for encouraging participation in incident reporting systems. There are other factors that contribute to such a reporting culture.

- *local champions.* ‘Local champions’ promote the system and who act as guarantors. They ensure that assurances of anonymity will be preserved in the face of external or managerial pressures. The previous chapters have already cited the role of David Wright in the local clinical system within Edinburgh’s Western General hospital. However, similar comments can be made about some of the much larger systems that operate within major companies. For instance, Capt. Mike O’Leary performs a similar function within British Airways’ confidential human factors reporting system [659]. A number of incident reporting systems have explicitly recognised the importance of these individuals. For instance, the Royal College of Anaesthetists advocates the identification of a Critical Incident coordinator who is responsible for drawing up and monitoring the operation of the system [715]. The explicit identification of an individual coordinator is a deliberate policy which goes beyond the more usual use of a committee structure within UK healthcare. There is, of course, a danger that the removal of such key individuals will threaten employee confidence and through that may jeopardise the continued operation of the system.

- *publicised participation.* One means of encouraging participation is to publish information about contribution rates from different groups within the organisation. This can illustrate that others have confidence in the system. However, this approach requires careful planning if it is not to have the opposite effect. In particular, it can be counter-productive to insist on reporting quotas. This can lead to fundamental questions about the purpose of a system that expects a certain number of failure reports from its staff within a particular interval. This is illustrated by a quotation from British Energy’s annual report on safety performance:

  > “The reportable events indicator is a measure of safety performance but more important than the number itself is the severity of the events reported. No target is set for this indicator in case this should discourage reporting. Indeed, within a healthy safety culture, the introduction of a ‘blame free’ reporting system may well cause an increase in the number of events reported.” [707]

  It is important that employees are provided with information about the number of contributions as they provide an important indicator of the health of the system. This quotation is instructive in this respect because it clearly links a blame free environment, a healthy safety culture and the elicitation of increasing numbers of incident reports.

- *maintaining the employees' voice.* One of the key elements in establishing what we have termed a 'reporting culture' is the preservation of the employees' voice from the moment at which an incident is identified to the final publication of feedback reports. There is often a danger that the employee's intentions in submitting a report will be turned to suit existing management priorities. Alternatively, genuine concerns may be lost in the process of filtering that was described in the previous chapter. This is a non-trivial task, especially when a technical analysis is required to identify the underlying causes of an incident. For example, an incident investigator recently told me of how he had tried to explain that there were extenuating circumstances, including distractions and shift patterns, that had contributed to an 'error'. This individual refused to listen to these arguments; preferring to accept blame for the incident. They insisted that interpretation must be included in the final report. Such situations create considerable conflicts for analysts who want to retain the support of the contributors while at the same time provide an accurate overview of the causal factors that lead to an incident.

- *system visibility.* It is also important that potential contributors are made aware of the procedures and mechanisms that support an incident reporting system. They must know *how* to report an adverse occurrence or a safety concern. Other aspects of system visibility can contribute to a reporting culture. For instance, the system should receive adequate resourcing so that prompt feedback can be provided. This is critical in creating an impression that contributions will be taken seriously. Reporting systems should also be visible at a corporate level if employees are to be confident that their views will have a strategic effect, in addition to any short-term changes that might be instigated within a particular team or group.

The following quotation further illustrates how British Energy has promoted its reporting system. In contrast to the previous citation, this excerpt focuses on the safety systems that are in operation at one site, Hinkley Point B reactor, within the organisation. The reporting system is considered to be an integrated part of wider mechanisms that are designed to ensure employee safety and to protect the environment. The following quotation is particularly interesting because it explains how an observational monitoring system has not yet been implemented. Previous chapters in this book have already argued that such observational studies are necessary in order for analysts to assess the importance of particular incidents within the wider context of operator tasks. The reference to the system at the highest level within the organisations safety plan makes it visible and reinforced it's importance to workers, managers and regulators:

> Hinkley Point B's safety performance continued to improve and the station met seven of the eight targets it set. The ISRS level achieved was 7. A RoSPA Gold Award was received for the first time. A Safety Information Centre, for the use of everyone on the station site, has been set up. A contractors' Health and Safety Committee has encouraged development and sharing of best practices. Near-miss reporting has contributed to safety performance. The independently audited housekeeping score was better than that targeted. The number of outstanding safety modifications has been reduced below the target level set. The one target missed was the aim of introducing non-obtrusive behaviour monitoring, based on self-assessment. This target has been carried forward to next year.
>
> Safety Awareness Week laid on an impressive programme of events and exhibitions involving the local community, emergency services and contractors. Celebrity input came from Geoff Capes who, appropriately, demonstrated manual handling techniques. The station successfully reduced its collective radiation dose below target by improved working practices, despite two periods of man entry into the reactor pressure vessel, one unforeseen at the start of the year...
>
> ENVIRONMENTAL A Station Environmental Plan aids a commitment to continuous improvement under ISO 14001 to which the station successfully converted from BS 7750. The station met all of its environmental objectives. It reduced the quantity of LLW

> it produced and improved contingency plans for dealing with oil and chemical spills. Development of the station nature trail continued with habitat management and creation of a wildflower meadow. There are over 1,000 species of flowering plants and invertebrates on the trail, more than 100 of which are currently listed in the Somerset Wildlife Trust list of notable species... A Peregrine Falcon nest ledge on the reactor building has been added to the five other raptor nest boxes situated in and around the nature trail." (Location report: Hinkley point B, [707])

A continuing theme in the justifications that support many incident reporting systems is that they increase the visibility of potential failures to many different groups within a workforce. This creates a recursive argument. Reading about incidents can increase an individual's sensitivity to potential failures. They are more likely to notice other potential problems and this, in turn, may make them more likely to contribute reports to that same system. In other words, feedback about previous incidents is, arguably, the most important means of ensuring participation in a reporting system.

Providing Feedback

Effective invention to address acknowledged safety concerns provides what is arguably the most persuasive means of encouraging staff to participate in incident reporting schemes. At the highest level, feedback about safety improvements can be provided through staff publications that record the severity of incidents that are reported each year. For example, Table 5.2.2 presents incident statistics published by the UK Atomic Energy Authority [534]. It shows the number of incidents reported at each level of the International Atomic Energy Authority's International Nuclear Event Scale (INES). INES is used to provide means of comparison between the reports that different national systems submit to the IAEA's INES incident database. Incidents at level 1 are simply regarded as anomalys, level 2 is an incident, level 3 is a serious incident, level 4 is an accident with significant off-site risk, level 5 is an accident with off-site risk, level 6 is a serious accident, level 7 is a major accident. The Chernobyl was classified at level 7, while the 1989 incident at the Vandellos nuclear power plant in Spain was classified at level 3. This did not result in an external release of radioactivity, nor was there damage to the reactor core or contamination on site. Fire did, however, damage the plant's safety systems. The IAEA does not provide examples of incidents below level 3; this is left to the prerogative of individual states. The benefit of this style of feedback is that managers and operators can compare national or local safety standards against those of other countries. For example, in 1997 the total INES summary produced for the IAEA recorded 16 anomalies at level 1, 15 incidents at level 2, 2 serious incidents at level 3 and no accidents between levels 4 and 7.

Year	INES level 1	INES level 2	INES levels 3-7
1996/97	4	3	0
1997/98	1	0	0
1998/99	1	1	0

Table 5.2: UK AEA Incident Statistics 1996-1999

The data presented in Table 5.2.2 is at a very high level of abstraction. Individual workers must relate such high-level categorisations to the risks that they face in the everyday tasks. This is not straightforward and, indeed, it is questionable whether such statistics would ever have a direct effect on future contribution rates to incident reporting schemes. On the other hand, staff may also receive a far more detailed level of feedback about the ways in which particular sets of incident data have been used more directly to address common safety concerns in many different incident reports. For example, the following quotation comes from Boeing's Aero magazine. This publication often describes ways in which company personnel and Boeing/Douglas operators have used incident reports to provide insights into technical problems. It is important to note that quotation begins by stressing the role of incident reports within improved training material. It then goes on to identify this material as a key factor in the reduction of rejected takeoff incidents during the 1990s:

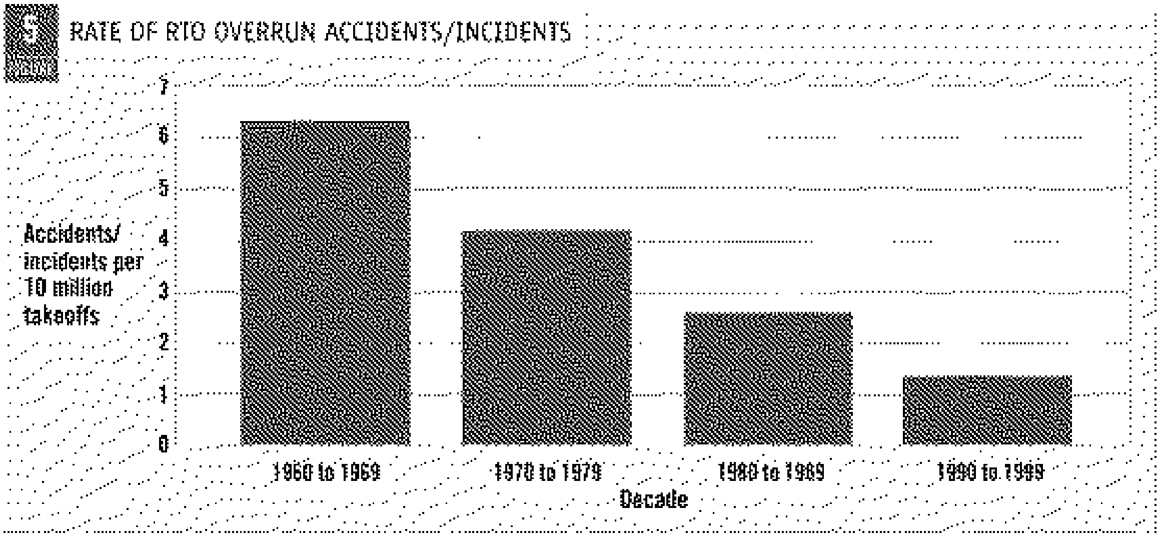

Figure 5.2: Accident and Incident Rates for Rejected Takeoff Overruns

> "The Takeoff Safety Training Aid (TOSTA) contains a list of the 74 Rejected Takeoff (RTO) overrun accidents and incidents studied during development of the training aid... Unfortunately, RTO overrun accidents and incidents continue to occur. However, the rate of occurrence continues to drop. Figure 5.2 (in this document) shows the rate of RTO overrun accidents and incidents expressed as events per 10 million takeoffs. Compared to the 1960s, the 1990s showed a 78 percent decrease in the rate of RTO overrun accidents and incidents. The industry can attribute this major improvement in RTO safety to many factors, but especially to better airplane systems, better and more reliable engines and in the 1990s, better training and standards."[510]

This quotation illustrates how statistical information about incidents and accidents can be used to provide feedback about the initiatives that are intended to avoid the recurrence of previous failures. This approach does, however, raise a number of important questions about the role of statistical feedback in encouraging participation in incident reporting systems:

- *too abstract and difficult to relate to everyday tasks.* As mentioned above, it can be difficult to map from high-level statistics down to the daily safety concerns that often persuade individual's to contribute to reporting systems. In particular, high level categorisations provide little or no information about the sorts of incidents that fall within the scope of the system. Finally, it can be difficult for individual's to determine whether others within their working teams or local organisations are also participating in the systems.

- *the paradox of low numbers may dissuade further participation.* This paradox centres on the idea that workers can be dissuaded from contributing reports if they see that only a few submissions are ever made. as mentioned in previous chapters, there is a very real concern that individuals may be identified and singled-out as trouble makers. In consequence, a high level of contribution at a low level of criticality can be taken to provide an indication of a positive safety culture. However, much of the statistical feedback provided to users often focuses on reductions in the already small number of high-criticality events.

- *they focus on structural problems that individuals cannot effect.* Regulators and safety managers must, typically, monitor incident data. They must ensure that any 'statistically significant' incidents are addressed through necessary investment, including improved operator training. As a result, statistical summaries often provide insights into problems that have already been solved or about issues that lie beyond the immediate influence of those who contribute to a reporting system.

- *too much emphasis on solved problems.* Statistical summaries are often used to evaluate the effectiveness of incident reporting systems. These summaries can then be used to encourage future submissions. However, as mentioned above, this need to encourage participation can also have undesired side-effects. In particular, the publication of data about previous successes can persuade operators that the base safety level of an application has been raised to a point where it is no longer necessary to report particular occurrences. Earlier sections have, however, pointed out that some systems, such as TCAS, have reduced certain froms of incident but have also contributed to other new adverse occurrences. Publishing 'raw' data about reduced proximity violations through the introduction of TCAS might help to obscure the continuing problems that these systems are posing for Air Traffic Management. There is a danger, therefore, that statistical summaries about the effectiveness of incident reporting can lead to undue complacency.

Many reporting systems avoid these criticisms by supplementing raw statistical information with more qualitative accounts and anecdotal editorials about previous incidents. For example, the image on the left of Figure 5.3 illustrates the Feedback reports that are produced by the Confidential Human Factors Incident Reporting Programme (CHIRP) . Feedback is distributed to personnel within commercial and general aviation. They are provided in paper form. They are also available on-line in HTML and PDF formats that can easily be both downloaded and printed. The CHIRP Feedback newsletter has a circulation of approximately 30,000. The ASRS's equivalent publication, Callback, is distributed to 85,000 aviation professionals. As can be seen from the cover in Figure 5.3 statistical data about the frequency and nature of submissions, typically in the form of pie-charts, introduces more qualitative accounts. These are intended to speak 'with the voice' of the individuals who are concerned in an incident:

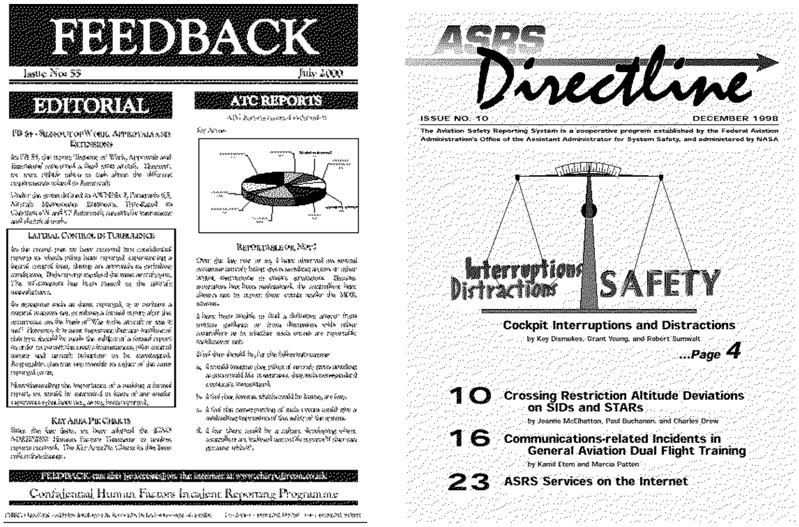

FEEDBACK

Issue No: 55 July 2000

EDITORIAL

[illegible]

LATERAL CONTROL IN TURBULENCE

[illegible]

KEY AREA PIE CHARTS

[illegible]

ATC REPORTS

[illegible]

FEEDBACK can also be accessed on the internet at: [illegible]

Confidential Human Factors Incident Reporting Programme

ASRS Directline

ISSUE NO. 10 DECEMBER 1998

The Aviation Safety Reporting System is a cooperative program established by the Federal Aviation Administration's Office of the Assistant Administrator for System Safety, and administered by NASA

Cockpit Interruptions and Distractions
by Key Dismukes, Grant Young, and Robert Sumwalt
...Page 4

10 Crossing Restriction Altitude Deviations on SIDs and STARs
by Jeanne McElhatton, Paul Buchanan, and Charles Drew

16 Communications-related Incidents in General Aviation Dual Flight Training
by Kamil Etem and Marcia Patten

23 ASRS Services on the Internet

Figure 5.3: CHIRP and ASRS Publications

> "Repetitive Defect and Sign-offs
>
> Yet another example of why maintenance engineering management should not be allowed to hold certifying approvals. The aircraft had several occurrences of No 1 engine fire detection loop failure on test. The usual steps were taken by line personnel (connectors cleaned) etc. up to AND including replacing the fire loop. As the defect was

> intermittent, it slipped through and reared its ugly head again the next day during crew checks. It finally reached the point where the line avionics personnel refused to 'shake it up' to get it going, the system needed proper down-time for investigation. Yet on four continuous reports, an A and C engineer with NO avionics clearance or know-how, released the aircraft to service with an inoperative fire detection system. This engineer was a mid-level manager with both a cavalier attitude to anything non-mechanical and also under pressure from management above him. What steps are being taken to address management's limitations to release aircraft to service?
>
> *Editorial comment: The alleged circumstances relating to the release of the aircraft were investigated by CAA (SRG) and corrective actions agreed. In the case of a repetitive defect that has not been cleared after three attempts, the procedure requires that the aircraft be withdrawn from service until the defect is rectified.*" [177]

This extract illustrates the way in which Feedback uses the contributor's own words to introduce safety concerns. This is direct and highly effective approach is also exploited by the ASRS' Callback publication. As in the previous example, editorial comments are used sparingly to indicate links with previous incidents, to point to corrective actions that personnel can take and, as in this example, to follow-up actions that the reporting organisation have instigated in response to the contribution. This final point is particularly important; it confirms that actions can and will be taken when safety concerns are elicited by a reporting system.

The image on the right of Figure 5.3 illustrates a slightly different form of feedback from the CHIRP publication. The ASRS' DirectLine journal is intended to support operators and flight crews of commercial carriers and corporate fleets. Unlike Callback and Feedback there is a greater degree of editorial comment in this publication. The articles in DirectLine, typically, address a particular issue that has been raised in a number of different contributions. For instance, the previous reports about TCASII were all drawn from a DirectLine study about the use and ab-use of this system. The following excerpt illustrates the difference in tone between Callback/Feedback and DirectLine, it is drawn from a study on cockpit interruptions:

> "Why do activities as routine as conversation sometimes interfere with monitoring or controlling the aircraft? Cognitive research indicates that people are able to perform two tasks concurrently only in limited circumstances, even if they are skillful in performing each task separately. Broadly speaking, humans have two cognitive systems with which they perform tasks; one involves conscious control, the other is an automatic system that operates largely outside of conscious control. The conscious system is slow and effortful, and it basically performs one operation at a time, in sequence. Learning a new task typically requires conscious processing, which is why learning to drive a car or fly an airplane at first seems overwhelming: the multiple demands of the task exceed conscious capacity. Automated cognitive processes develop as we acquire skill; these processes are specific to each task, they operate rapidly and fluidly, and they require little effort or attention." [213]

As mentioned, the intentions behind DirectLine are quite different from those of its sister publications. One consequence of this is that it plays a different role in the elicitation of future contributions. One potential effect is that it sensitises others within the aviation community to the importance of particular incidents which are symptomatic of deeper underlying problems; such as cockpit distractions in the previous example. DirectLine also helps to demonstrate ways in which incident reporting can be integrated into wider safety concerns within the aviation industry. Rather than simply picking out individual incidents for editorial comment, this publication points to clusters of similar events. This, in turn, has had a considerable influence on developers, operators and regulators. A point that is illustrated by the previous quotation from Boeing's Aero article on Rejected Takeoff (RTO) overrun accidents and incidents.

Callback, Feedback and DirectLine help to elicit further contributions by explaining how previous incidents can be avoided in the future. These publications are all accessible in electronic form, over the Internet. However, they all rely upon a traditional format. These publications exploit the linear style of conventional newsletters or journals. This has important benefits. In particular, they can be easily printed for wider dissemination. However, a number of incident investigation authorities are exploiting alternative approaches. Most of these are based around providing Internet access to databases of previous incidents. This approach is partly exploited by the ASRS . Reports are published incrementally so that the fifty most recent contributions are summarised in each batch. However, other organisations extend this database approach to include not simply summaries of the incident but also information about the associated investigation and analysis. For instance, the interface on the left of Figure 5.4 provides access to the US Chemical Safety and Hazard Investigation Board's incident reports 'Centre' [162]. This provides access to reports on both accidents, involving fatalities, and incidents. Users can search through an archive of incident reports using a number of different tools. The image on the right of Figure 5.4 illustrates the information that is returned about each incident. There are a number of innovative features about this application that encourage contributions about adverse occurrences. In particular, it is possible to chart the course of an investigation as it progresses. This provides individual contributors with confirmation that their reports are being acted upon. Search facilities also enable potential contributors to determine whether other similar occurrences have been notified to the system. Chapter 14 will provide a more detailed analysis of the dissemination techniques that can be recruited both to publicise the findings of incident investigations and to elicit further contributions to reporting systems. In contrast, this section continues to examine other means of encouraging the manual submission of information about adverse occurrences.

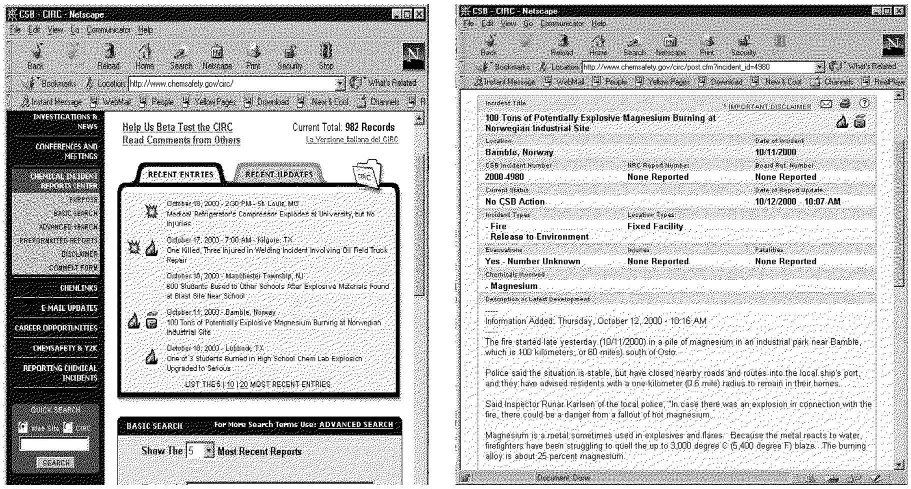
CSB - CIRC - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://www.chemsafety.gov/circ/ What's Related
Instant Message WebMail People Yellow Pages Download New & Cool Channels
INVESTIGATIONS & NEWS
CONFERENCES AND MEETINGS
CHEMICAL INCIDENT REPORTS CENTER
PURPOSE
BASIC SEARCH
ADVANCED SEARCH
PREFORMATTED REPORTS
DISCLAIMER
COMMENT FORM
CHEMLINKS
E-MAIL UPDATES
CAREER OPPORTUNITIES
CHEMSAFETY & Y2K
REPORTING CHEMICAL INCIDENTS
QUICK SEARCH
Web Site CIRC
SEARCH
RECENT ENTRIES
RECENT UPDATES
CIRC
BASIC SEARCH
For More Search Terms Use: ADVANCED SEARCH
CSB - CIRC - Netscape
Location: http://www.chemsafety.gov/circ/post.cfm?incident_id=4980
RealPlayer
IMPORTANT DISCLAIMER
Document: Done

Help Us Beta Test the CIRC

Read Comments from Others

Current Total: **982 Records**

La Versione Italiana del CIRC

October 18, 2000 - 2:30 PM - St. Louis, MO
Medical Refrigerator's Compressor Explodes at University, but No Injuries

October 17, 2000 - 7:00 AM - Kilgore, TX
One Killed, Three Injured in Welding Incident Involving Oil Field Truck Repair

October 16, 2000 - Manchester Township, NJ
600 Students Bused to Other Schools After Explosive Materials Found at Blast Site Near School

October 11, 2000 - Bamble, Norway
100 Tons of Potentially Explosive Magnesium Burning at Norwegian Industrial Site

October 10, 2000 - Lubbock, TX
One of 3 Students Burned in High School Chem Lab Explosion Upgraded to Serious

LIST THE 5 | 10 | 20 MOST RECENT ENTRIES

Show The 5 Most Recent Reports

Incident Title
100 Tons of Potentially Explosive Magnesium Burning at Norwegian Industrial Site

Location		Date of Incident
Bamble, Norway		**10/11/2000**
CSB Incident Number	NRC Report Number	Board Ref. Number
2000-4980	**None Reported**	**None Reported**
Current Status		Date of Report Update
No CSB Action		**10/12/2000 - 10:07 AM**
Incident Types	Location Types	
- Fire **- Release to Environment**	**Fixed Facility**	
Evacuations	Injuries	Fatalities
Yes - Number Unknown	**None Reported**	**None Reported**
Chemicals Involved		
- Magnesium		

Description or Latest Development

Information Added: Thursday, October 12, 2000 - 10:16 AM

The fire started late yesterday (10/11/2000) in a pile of magnesium in an industrial park near Bamble, which is 100 kilometers, or 60 miles) south of Oslo.

Police said the situation is stable, but have closed nearby roads and routes into the local ship's port, and they have advised residents with a one-kilometer (0.6 mile) radius to remain in their homes.

Said Inspector Runar Karlsen of the local police, "In case there was an explosion in connection with the fire, there could be a danger from a fallout of hot magnesium.

Magnesium is a metal sometimes used in explosives and flares. Because the metal reacts to water, firefighters have been struggling to quell the up to 3,000 degree C (5,400 degree F) blaze. The burning alloy is about 25 percent magnesium.

Figure 5.4: Web Interface to the CHSIB Incident Collection

Publicising Procedures and Scoping the System

Chapter 1 introduced some of the problems that arise when attempting to define what are, and what are not, abnormal occurrences. This is not simply a research issue; it is of fundamental importance for individuals who must determine whether or not an incident is worth reporting. Exhaustive, or closed, definitions rely upon pre-defined lists of abnormal events. There are very few examples of such systems because they, typically, place undue constraints on what should be reported. Closed definitions dissuade individuals from contributing relevant information about other types of incidents. This is a particular problem if new technology or working practices leads to different types of occurrences that do not appear on the list.

Alternatively, open approaches provide broad definitions of what are critical incidents. They are, however, open to subjective biases. Different individuals have very different opinions about what should be reported. For example, the Royal College of Anaesthetists incident reporting form contains the following definition:

> A Critical Incident may be defined as an event which led to harm, or could have led to harm, if it had been allowed to progress. It should be preventable by a change of practice. Complications are occurrences of patient harm, and are sometimes the outcome of critical incidents. If you experience what you think is a critical incident whether or not it has such an adverse outcome, please fill in this form as soon as possible after the event - memory changes very rapidly." [715]

Inductive guidelines provide more limited examples of critical incidents than the exhaustive approach mentioned above. Pragmatically, most systems exploit a combination of open definitions and inductive guidelines. For instance, the Royal College's form provides examples of possible incidents when it considers the different levels of 'preventability' that might associated with an occurrence. This guidance is important because it implicitly also provides information about what events are considered to be within the scope of the system. This may guide the elicitation of reports within these categories:

> "Please grade how PREVENTABLE the incident or complication was as follows:
>
> 1. Probably preventable within current resource (e.g. failure to do preop machine check);
> 2. Probably preventable with reasonable extra resource (e.g. failure to detect oesophageal intubation would be improved by having capnographs);
> 3. Possibly preventable within current resource (e.g. pneumothorax during CVP insertion might be prevented by better teaching and supervision);
> 4. Possibly preventable with reasonable extra resource (e.g. problem arising because anaesthetist unwell might be prevented by more cover);
> 5. Not obviously preventable by any change in practice (e.g. electricity grid failure)" [715]

However, there are other alternatives. For example, the US Department of the Energy's Computerised Accident/Incident Reporting System (CAIRS) uses a definition based on monetary loss:

> "The reporting criteria for CAIRS injury/illness cases changed, effective January 1, 1990, from the criteria specified in the DOE Guide to the Classification of Recordable Accidents to the Occupational Safety and Health Administration (OSHA) guidelines. The reporting threshold for property damage cases was originally set at $1,000 and remained that way until January 1, 1996, when it was raised to $5,000. The vehicle accident reporting threshold was $250 from 1975 through 1985, $500 from 1986 through 1995, and was raised to $1,000 effective January 1, 1996." [655]

The problem with this approach is that it can be difficult to anticipate the potential losses that might be experienced from near-miss incidents. As mentioned in previous chapters, there is also a danger of under-reporting if potential contributors under-assess the amount of damage caused by an incident for whatever reason. For any definition of an incident, there are a number of fundamental principles that must be followed:

- *it is important to publish guidance on the scope of reports.* This may seems obvious. However, it is critical that scope and type of occurrences are published. The fear of retribution or disclosure are powerful disincentives not to contribute if there is any doubt about whether or not an occurrence falls within the scope of the system. From this it follows that any definition must be clearly understood and accepted by potential contributors. Staff must be explicitly trained to use open definitions so that they can consistently identify those occurrences that should be reported. This is particularly important during the start-up phase of any system when potential contributors may not have the feedback reports that provide more detailed examples of what should be reported.

- *the scope of the system should conform to national and international standards.* As mentioned in previous chapters, there are numerous national and international guidelines which specify what must be reported within some systems. These guidelines are, typically, intended to ensure that different classes of events are treated in a consistent manner. This, in turn, enables information to be exchanged between different countries. In particular, the frequency of incidents at the same level of criticality is often used as a comparative measure of national safety performance. For instance, this is a primary motivation behind the International Atomic Energy Authority's International Nuclear Event Scale (INES) and the severity assessments in the International Civil Aviation Organisation's (ICAO) Annex 13 guidance.

- *allow for local differences if properly justified and documented.* Local circumstances also affect what is, and what is not, covered by incident reporting schemes. Regional managers often decide to introduce particular adverse occurrences into their system if they perceive that they pose a particular local risk. This might be done if those occurrences are abnormally frequent or if local conditions mean that those events carry unusually high consequences for their region. These regional differences must not jeopardise the minimum reporting criteria established by national and international systems. It is equally important that the scope of reporting systems can be informed by local experience. These local concerns must be explained to potential contributors if they are to guide the submission of incident reports.

- *it is important to monitor contributions and update definitions.* It has also been argued that closed lists and open definitions of adverse occurrences can dissuade potential contribtuons to incident reporting systems. In consequence, most systems exploit open definitions backed with a number of examples to illustrate what falls within the scope of the system. The success of this approach can be monitored by the range of contributions that are received. As mentioned above, special initiatives and tailored reporting forms can be used to address apparent omissions by focusing attention on particular types of occurrences. For instance, new installations or operating procedures, such as parallel approaches in Air Traffic Management, can encourage managers to re-iterate the importance of reporting even low criticality failures involving these new systems. Again, the practical implementation of these monitoring techniques creates particular problems during the start-up phase when there will be little or no baseline figures for comparison. This is a particular problem for systems that monitor for potential problems with new equipment; relatively few submissions may indicate a successful application or an unsuccessful reporting system! Baseline data can be obtained by analysing the frequency of trigger events recorded using automated monitoring equipment. Alternatively, observational studies can be used to provide more direct qualitative information that supplements the insights that are contributed through voluntary reporting systems.

Previous sections have argued that incident reporting systems depend upon the elicitation of information about potential failures or previous adverse occurrences. This, in turn, depends upon the successful design of incident reporting forms. Poorly constructed forms can lead to confusion about the information that is being requested. Such assessments must, however, be balanced by the observation that relatively little is known about the impact of form design upon reporting behaviour. The following paragraphs, therefore, use a comparative study of existing incident forms to identify key decisions that must be made during the design of future documents that elicit reports about adverse occurrences.

5.3 Form Contents

Hundreds of local, national and international systems are using ad hoc, trial and error techniques to arrive at the appropriate content and layout of the forms that are used to elicit incident reports. As a result, there is a huge variation in both the information that is requested from the user and the information that is provided to prompt them for relevant information. For example, some schemes have found it useful to print the forms that elicit future submissions on the back of the newletters and journals that publicise information about previous incidents. Other systems rely entirely on

Internet-based electronic forms. In spite of this diversity, it is possible to identify a number of common features that are shared by many reporting systems. For instance, Table 5.3 summarises the information that is typically requested by these forms. As we shall see, it is not simply enough to request information about the incident itself. It is also important to identify ways in which safety systems successfully intervened to detect and to mitigate the consequences of an adverse occurrence. The following sections look beyond this high level classification to look at the different techniques that have been exploited by a number of existing local and national systems.

5.3.1 Sample Incident Reporting Forms

As mentioned, there are several different approaches to the presentation and dissemination of incident reporting forms. For example, some organisations provide printed forms that are readily at hand for the individuals that work within particular environments. This approach clearly relies upon the active monitoring of staff who must replenish the forms and who must collected completed reports. The form shown in Figure 5.5 illustrates this approach.

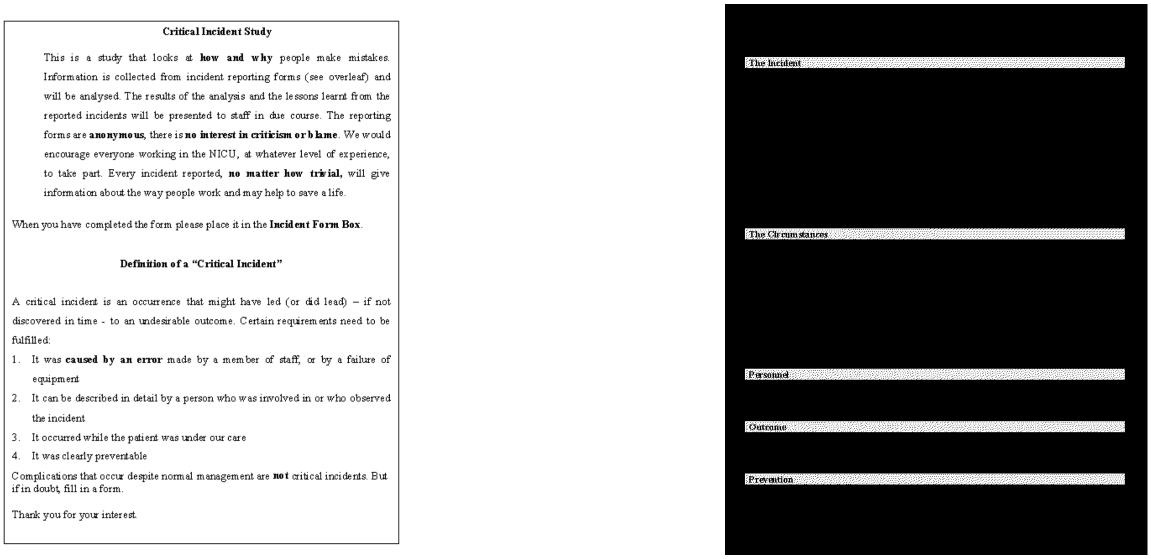

Critical Incident Study

This is a study that looks at **how and why** people make mistakes. Information is collected from incident reporting forms (see overleaf) and will be analysed. The results of the analysis and the lessons learnt from the reported incidents will be presented to staff in due course. The reporting forms are **anonymous**, there is **no interest in criticism or blame**. We would encourage everyone working in the NICU, at whatever level of experience, to take part. Every incident reported, **no matter how trivial,** will give information about the way people work and may help to save a life.

When you have completed the form please place it in the **Incident Form Box**.

Definition of a "Critical Incident"

A critical incident is an occurrence that might have led (or did lead) – if not discovered in time - to an undesirable outcome. Certain requirements need to be fulfilled:

1. It was **caused by an error** made by a member of staff, or by a failure of equipment
2. It can be described in detail by a person who was involved in or who observed the incident
3. It occurred while the patient was under our care
4. It was clearly preventable

Complications that occur despite normal management are **not** critical incidents. But if in doubt, fill in a form.

Thank you for your interest.

The Incident

The Circumstances

Personnel

Outcome

Prevention

Figure 5.5: Incident Reporting Form for a UK Neonatal Intensive Care Unit [119]

The document in Figure 5.5 was developed for a Neonatal Intensive Care Unit and is based upon a form that has been used for almost a decade in an adult intensive care environment [122, 119]. As can be seen, this form relies upon free-text fields where the user can describe the incident that they have witnessed. This approach works because the people analysing the report are very familiar with the Units that exploit them. In contrast, national and international schemes typically force their respondents to select their responses from lists of more highly constrained alternatives. For example, NASA and the FAA's ASRS scheme covers many diverse occupations, ranging from maintenance to aircrew activities, in the many different geographical regions of the United States. This has a radical effect on forms such as that shown in Figure 5.6 which is designed to elicit reports about Air Traffic Control incidents. Pre-defined terms are used to distinguish between the many different control positions and activities that are involved in an international, air traffic control system. Much of this activity information remains implicit in local forms such as that shown in Figure 5.5.

The local reporting form shown in Figure 5.5 is distributed by placing paper copies within the users' working environment. In contrast, ASRS forms are also available over the World Wide Web . They can be downloaded and printed using Adobe's proprietary Portable Display Format (PDF). The geographical and the occupational coverage of the ASRS system again determine this approach. The web is perceived to provide a cost-effective means of disseminating incident reporting form.

Topic of question:	Examples of information requested
Identification information:	Name, working team or unit, control centre information, current status of license.
Shift information:	When did the occurrence occur? When was their last break and for how long was it? When did they last operate this shift pattern in this control position? Were you training (or being trained?).
Station configuration:	What was the station configuration/manning like at the time of the occurrence? What was the ATC display configuration? Were you working with headsets/telephones/microphone and speaker? Were there any technical failures?
Operating characteristics:	What was the traffic volume like in your estimation? What was your workload like immediately before the occurrence? Were there any significant meteorological conditions?
Detection and mitigation factors:	What made you aware of the occurrence (e.g. automated warning, visual observation of radar)? Were there any circumstances that helped to mitigate any potential impact of the occurrence?
Other factors:	Are there any personal (off the job) circumstances that might affect the performance of you or others during the occurrence?
Free-text description of the occurrence:	Describe the occurrence and your performance/role during it. Also consider any ways in which you think that the occurrence might have been avoided.

Table 5.3: Developing Reporting Forms

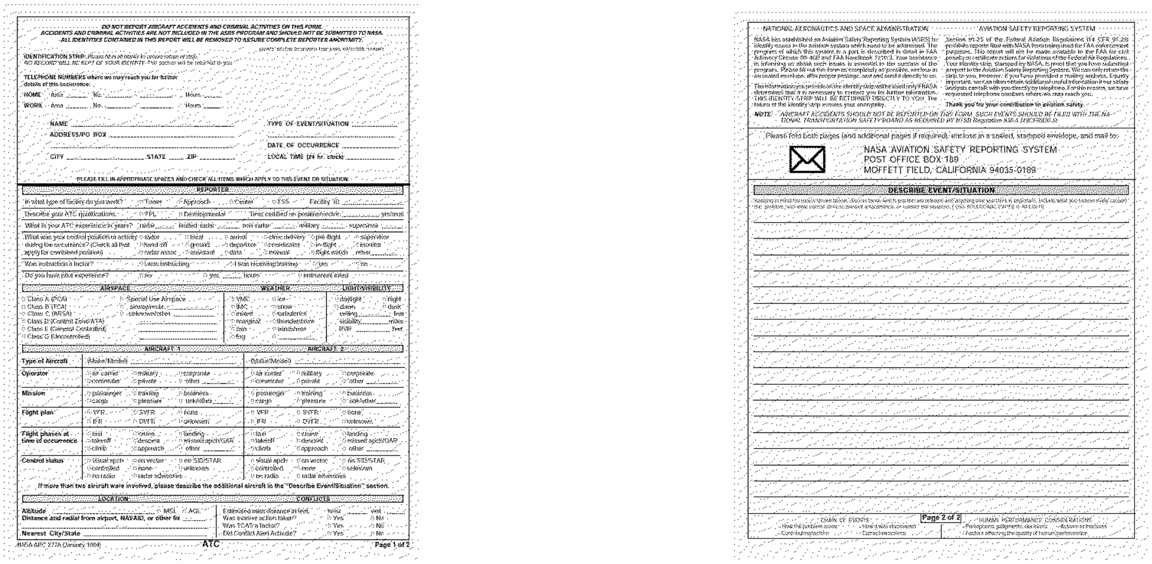

DO NOT REPORT AIRCRAFT ACCIDENTS AND CRIMINAL ACTIVITIES ON THIS FORM.
ACCIDENTS AND CRIMINAL ACTIVITIES ARE NOT INCLUDED IN THE ASRS PROGRAM AND SHOULD NOT BE SUBMITTED TO NASA.
ALL IDENTITIES CONTAINED IN THIS REPORT WILL BE REMOVED TO ASSURE COMPLETE REPORTER ANONYMITY.

(SPACE BELOW RESERVED FOR ASRS DATE/TIME STAMP)

IDENTIFICATION STRIP: Please fill in all blanks to ensure return of strip.
NO RECORD WILL BE KEPT OF YOUR IDENTITY. This section will be returned to you.

TELEPHONE NUMBERS where we may reach you for further details of this occurrence:
HOME Area ______ No. ________ - ________ Hours ______
WORK Area ______ No. ________ - ________ Hours ______

NAME ______________________ TYPE OF EVENT/SITUATION ______________
ADDRESS/PO BOX ______________________
______________________ DATE OF OCCURRENCE ______________
CITY ____________ STATE ____ ZIP ________ LOCAL TIME (24 hr. clock) ______________

PLEASE FILL IN APPROPRIATE SPACES AND CHECK ALL ITEMS WHICH APPLY TO THIS EVENT OR SITUATION.

REPORTER

In what type of facility do you work? ○ Tower ○ Approach ○ Center ○ FSS Facility ID ______
Describe your ATC qualifications. ○ FPL ○ Developmental Time certified on position/sector: ______ years
What is your ATC experience in years? radar ______ limited radar ______ non-radar ______ military ______ supervisor ______
What was your control position or activity during the occurrence? (Check all that apply for combined positions) ○ sector ○ local ○ arrival ○ clearance delivery ○ pre-flight ○ supervisor ○ hand-off ○ ground ○ departure ○ coordinator ○ in-flight ○ monitor ○ radar assist ○ assistant ○ data ○ manual ○ flight watch other ______
Was instruction a factor? ○ I was instructing ○ I was receiving training ○ yes ○ no
Do you have pilot experience? ○ no ○ yes ______ hours ○ instrument rated

AIRSPACE		WEATHER		LIGHT/VISIBILITY	
○ Class A (PCA)	○ Special Use Airspace	○ VMC	○ ice	○ daylight	○ night
○ Class B (TCA)	○ airway/route ______	○ IMC	○ snow	○ dawn	○ dusk
○ Class C (ARSA)	○ unknown/other ______	○ mixed	○ turbulence	ceiling ______ feet	
○ Class D (Control Zone/ATA)	______	○ marginal	○ thunderstorm	visibility ______ miles	
○ Class E (General Controlled)	______	○ rain	○ windshear	RVR ______ feet	
○ Class G (Uncontrolled)	______	○ fog	○ ______		

	AIRCRAFT 1	AIRCRAFT 2
Type of Aircraft	(Make/Model) ______	(Make/Model) ______
Operator	○ air carrier ○ military ○ corporate ○ commuter ○ private ○ other ______	○ air carrier ○ military ○ corporate ○ commuter ○ private ○ other ______
Mission	○ passenger ○ training ○ business ○ cargo ○ pleasure ○ unk/other ______	○ passenger ○ training ○ business ○ cargo ○ pleasure ○ unk/other ______
Flight plan	○ VFR ○ SVFR ○ none ○ IFR ○ DVFR ○ unknown	○ VFR ○ SVFR ○ none ○ IFR ○ DVFR ○ unknown
Flight phases at time of occurrence	○ taxi ○ cruise ○ landing ○ takeoff ○ descent ○ missed apch/GAR ○ climb ○ approach ○ other ______	○ taxi ○ cruise ○ landing ○ takeoff ○ descent ○ missed apch/GAR ○ climb ○ approach ○ other ______
Control status	○ visual apch ○ on vector ○ on SID/STAR ○ controlled ○ none ○ unknown ○ no radio ○ radar advisories	○ visual apch ○ on vector ○ on SID/STAR ○ controlled ○ none ○ unknown ○ no radio ○ radar advisories

If more than two aircraft were involved, please describe the additional aircraft in the "Describe Event/Situation" section.

LOCATION	CONFLICTS
Altitude ______ ○ MSL ○ AGL	Estimated miss distance in feet: horiz ______ vert ______
Distance and radial from airport, NAVAID, or other fix ______	Was evasive action taken? ○ Yes ○ No
______	Was TCAS a factor? ○ Yes ○ No
Nearest City/State ______	Did Conflict Alert Activate? ○ Yes ○ No

NASA ARC 277A (January 1998) ATC Page 1 of 2

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION AVIATION SAFETY REPORTING SYSTEM

NASA has established an Aviation Safety Reporting System (ASRS) to identify issues in the aviation system which need to be addressed. The program of which this system is a part is described in detail in FAA Advisory Circular 00-46D and FAA Handbook 7210.3. Your assistance in informing us about such issues is essential to the success of the program. Please fill out this form as completely as possible, enclose in a sealed envelope, affix proper postage, and send it directly to us.

The information you provide on the identity strip will be used only if NASA determines that it is necessary to contact you for further information. THIS IDENTITY STRIP WILL BE RETURNED DIRECTLY TO YOU. The return of the identity strip assures your anonymity.

Section 91.25 of the Federal Aviation Regulations (14 CFR 91.25) prohibits reports filed with NASA from being used for FAA enforcement purposes. This report will not be made available to the FAA for civil penalty or certificate actions for violations of the Federal Air Regulations. Your identity strip, stamped by NASA, is proof that you have submitted a report to the Aviation Safety Reporting System. We can only return this strip to you, however, if you have provided a mailing address. Equally important, we can often obtain additional useful information if our safety analysts can talk with you directly by telephone. For this reason, we have requested telephone numbers where we may reach you.

Thank you for your contribution to aviation safety.

NOTE: AIRCRAFT ACCIDENTS SHOULD NOT BE REPORTED ON THIS FORM. SUCH EVENTS SHOULD BE FILED WITH THE NATIONAL TRANSPORTATION SAFETY BOARD AS REQUIRED BY NTSB Regulation 830.5 (49CFR830.5).

Please fold both pages (and additional pages if required), enclose in a sealed, stamped envelope, and mail to:

NASA AVIATION SAFETY REPORTING SYSTEM
POST OFFICE BOX 189
MOFFETT FIELD, CALIFORNIA 94035-0189

DESCRIBE EVENT/SITUATION

Keeping in mind the topics shown below, discuss those which you feel are relevant and anything else you think is important. Include what you believe really caused the problem, and what can be done to prevent a recurrence, or correct the situation. (USE ADDITIONAL PAPER IF NEEDED)

Page 2 of 2

CHAIN OF EVENTS
• How the problem arose • How it was discovered
• Contributing factors • Corrective actions

HUMAN PERFORMANCE CONSIDERATIONS
• Perceptions, judgments, decisions • Actions or inactions
• Factors affecting the quality of human performance

Figure 5.6: ASRS Reporting Form for Air Traffic Control Incidents (January 2000)

ASRS report forms cannot, however, be submitted using Internet based technology. There are clear problems associated with the validation of such electronic submissions. A small but increasing number of reporting systems have taken this additional step towards the use of the Web as a means both of disseminating and submitting incident reporting forms. For instance, Figure 5.7 illustrates part of the on-line system that has been developed to support incident reporting within Swiss Departments of Anaesthesia [756].

As with the ASRS system and the local scheme, the CIRS reporting form also embodies a number of assumptions about the individuals who are likely to use the form. Perhaps the most obvious is that they must be computer literate and be able to use the diverse range of dialogue styles that are exploited by the system. They must also be able to translate between the incident that they have witnessed and the various strongly typed categories that are supported by the form. For instance, users must select from one of sixteen different types of surgical procedure that are recognised by the system. Perhaps more contentiously they must also characterise human performance along eight Likert scales that are used to assess lack of sleep, amount of work-related stress, amount of non-work related stress, effects of ill or healthy staff, adequate or inadequate knowledge of the situation, appropriate skills and appropriate experience. The introspective ability to independently assess such factors and provide reliable self-reports again illustrates how many incidents reporting forms reflect the designers' assumptions about the knowledge, training and expertise of the target workforce.

5.3.2 Providing Information to the Respondents

The previous section has illustrated a number of different approaches to the elicitation of information about human 'error' and systems 'failure'. However, these different approaches all address a number of common problems. The first is how to address the problem of under-reporting discussed in the first half of this chapter? Incident reporting forms must encourage people to contribute information about the incidents that they observe.

Assurances of Anonymity or Confidentiality

Previous chapters have explored the consequences of operating either an anonymous or a confidential system. Each of the systems presented in the Section 2 illustrates a different approach to this issue. For example, NASA administers the ASRS on behalf of the FAA. They act as an independent agency

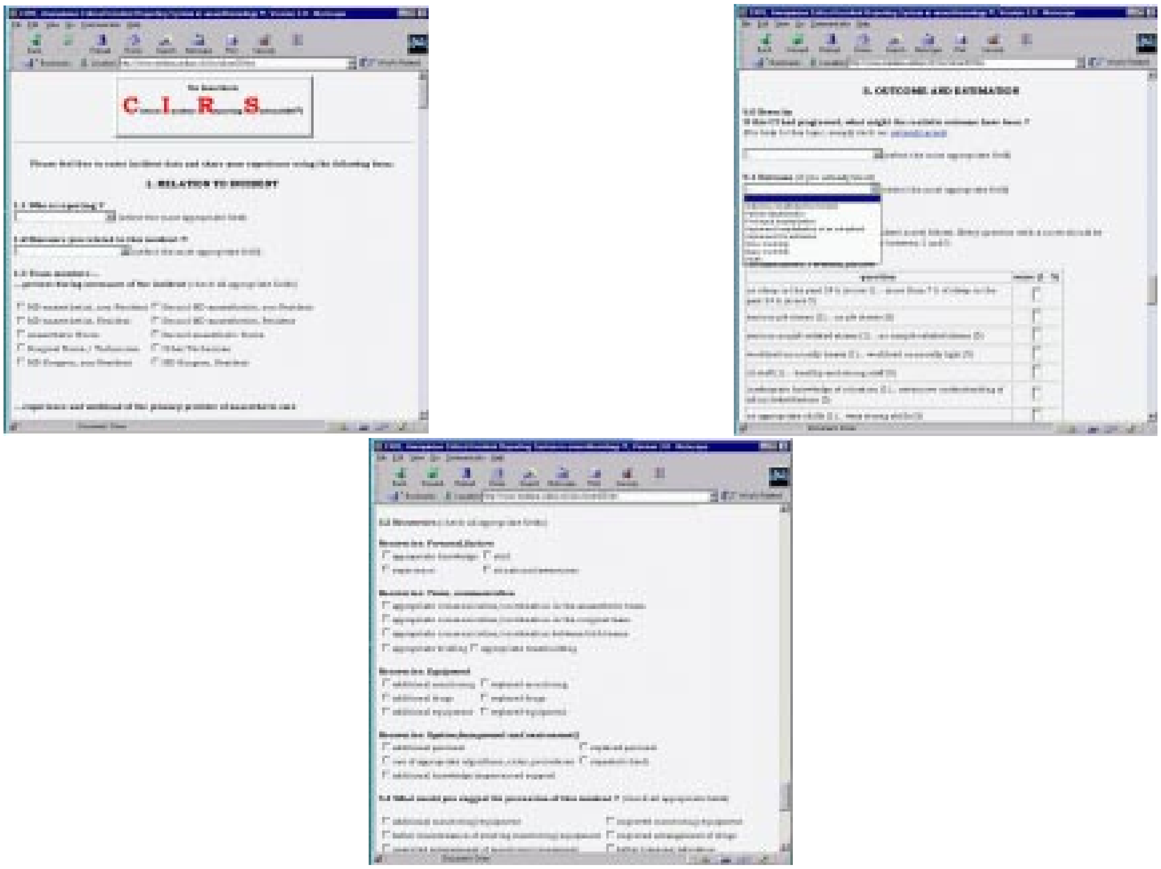

Figure 5.7: The CIRS Reporting System [756]

that guarantees the anonymity of respondents. FAA Advisory Circular Advisory Circular 00-46D states that:

> "The FAA will not seek, and NASA will not release or make available to the FAA, any report filed with NASA under the ASRS or any other information that might reveal the identity of any party involved in an occurrence or incident reported under the ASRS".

As mentioned, however, this scheme is confidential in the sense that NASA will only guarantee anonymity if the incident has no criminal implications. Respondents to the ASRS are asked to provide contact information so that NASA can pursue any additional information that might be needed. Conversely, the local scheme illustrated in Figure 5.5 does not request identification information from respondents. This anonymity is intended to protect staff and encourage their participation. However, it clearly creates problems during any subsequent causal analysis for reports of human error. It can be difficult to identify the circumstances leading to an incident if analysts cannot interview the person making the report.

However, this limitation is subject to a number of important caveats that affect the day to day operation of many local reporting schemes. For instance, given the shift system that operates in many industries and the limited number of personnel who are in a position to observe particular failures it is often possible for local analysts to make inferences about the people involved in particular situations. Clearly there is a strong conflict between the desire to prevent future incidents by breaking anonymity to ask supplementary questions and the desire to incidents by breaking anonymity to ask supplementary questions and the desire to safeguard the long-term participation of staff within the system. The move from paper-based schemes to electronic systems raises a host of complex social and technological issues surrounding the anonymity of respondents and the validation of submissions.

The Swiss scheme shown in Figure 5.7 states that:

> "During your posting of a case there will be NO questions that would allow an identification of the reporter, the patient or the institution. Furthermore we will NOT save any technical data on the individual reports: no E-mail address and no IP-number (a number that accompanies each submitted document on the net). So no unauthorised 'visitor' will find any information that would allow an identification of you or your patient or your institution (not even on our local network-computers) by browsing through the cases."

This addresses the concern that it is entirely possible for web servers to record the address of the machine making the submission without the respondent's knowledge. However, there is also a concern that groups might deliberately distort the findings of a system by generating spurious reports. These could, potentially, implicate third parties. It, therefore, seems likely that future electronic systems will follow the ASRS approach of confidential rather than anonymous reporting.

Definitions of an Incident?

It is important to provide users with a clear idea of when they should consider making a submission to the system. For example, the local scheme in Figure 5.5 states that an incident must fulfill the following criteria:

> "1. It was caused by an error made by a member of staff, or by a failure of equipment.
> 2. A person who was involved in or who observed the incident can describe it in detail.
> 3. It occurred while the patient was under our care. 4. It was clearly preventable.
>
> Complications that occur despite normal management are not critical incidents. But if in doubt, fill in a form."

This pragmatic definition from a long-running and successful scheme is full of interest for researchers working in the area of human error. For instance, incidents, which occur in spite of normal management, do not fall within the scope of the system. Some might argue that this effectively prevents the system from targeting problems within the existing management system. However, such criticisms neglect the focused nature of this local system, which is specifically intended to "target the doable" rather than capture all possible incidents.

In contrast to the local definition which reflects the working context of the unit in which it was applied, the wider scope of the CIRS approach leads to a much broader definition of the incidents under consideration:

> "Defining critical incidents unfortunately is not straightforward. Nevertheless we want to invite you to report your critical incidents if they match with this definition: an event under anaesthetic care which had the potential to lead to an undesirable outcome if left to progress. Please also consider any team performance critical incidents, regardless of how minimal they seem."

It is worth considering the implications of this definition in the light of previous research in the field of human error. For example, Reason has argued that many operators spend considerable amounts of time interacting in what might be terms a 'sub-optimal' manner [699]. Much of this behaviour could, if left unchecked, result in an undesirable outcome. However, operating practices and procedures help to ensure safe and successful operation. From this it follows that if respondents followed the literal interpretation of the CIRS definition then there could be a very high number of submissions. Some schemes take this broader approach one step further by requiring that operators complete an incident form after every procedure even if they only indicate that there had been no incident. The second interesting area in the CIRS definition is the focus on team working. The number of submissions to a reporting system is likely to fall as the initial enthusiasm declines. One means of countering this is to launch special reporting initiatives. For instance, by encouraging users to submit reports on particular issues such as team co-ordination problems. There is, however, the danger that this will lead to spurious attention being placed on relatively unimportant issues if the initiatives are not well considered.

The ASRS forms no longer contain an explicit indication of what should be reported. Paradoxically, the forms contain information about what is NOT regarded as being within the scope of the scheme.

> "Do not report aircraft accidents and criminal activities on this form".

This lack of an explicit definition of an incident reflects the success of the ASRS approach. In particular, it reflects the effectiveness of the feedback that is provided from the FAA and NASA. Operators can infer the sorts of incidents that are covered by the ASRS because they are likely to have read about previous incidents in publications such as the Callback magazine. This is distributed to more than 85,000 pilots, air traffic controllers and others in the aviation industry. Callback contains excerpts from ASRS incident reports as well as summaries of ASRS research studies. This coverage helps to provide examples of previous reporting behaviour. Of course, it might also be argued that apparently low participation rates, for example amongst cabin staff, could be accounted for by their relatively limited exposure to these feedback mechanisms. This raises further complications. In order to validate such hypotheses it is necessary to define an anticipated reporting rate for particular occupational groupings, such as cabin staff. This is impossible to do because without a precise definition of what an incident actually is, it is impossible to estimate exposure rates.

Explanations of Feedback and Analysis

Potential contributors must be convinced that their reports will be acted upon. For example, in the local system in Figure 5.5 includes the promise that:

> "Information is collected from incident reporting forms (see overleaf) and will be analysed. The results of the analysis and the lessons learned from the reported incidents will be presented to staff in due course."

This informal process is again typical of systems in which the lessons from previous incidents can be fed-back through ad hoc notices, reminders and periodic training sessions. It contrasts sharply with the ASRS approach:

> "Incident reports are read and analysed by ASRS's corps of aviation safety analysts. The analyst staff is composed entirely of experienced pilots and air-traffic controllers. Their years of experience are uniformly measured in decades, and cover the full spectrum of aviation activity: air carrier, military, and general aviation; Air Traffic Control in Towers, TRACONS, Centres, and Military Facilities. Each report received by the ASRS is read by a minimum of two analysts. Their first mission is to identify any aviation hazards, which are discussed in reports and flag that information for immediate action. When such hazards are identified, an alerting message is issued to the appropriate FAA office or aviation authority. Analysts' second mission is to classify reports and diagnose the causes underlying each reported event. Their observations, and the original de-identified report, are then incorporated into the ASRS's database."

The CIRS web-based system is slightly different from the other two examples. It is not intended to directly support intervention within particular working environments. Instead, the purpose is to record incidents so that other anaesthetists can access them and share experiences. It, therefore, follows that very little information is provided about the actions that will be taken in response to particular reports:

> "Based on the experiences from the Australian-Incident-Monitoring-Study, we would like to create an international forum where we collect and distribute critical incidents that happened in daily anaesthetic practice. This program not only allows the submission of critical incidents that happened at your place but also serves as a teaching instrument: share your experiences with us and have a look at the experiences of others by browsing through the cases. CIRS is anonymous."

This approach assumes that the participating group already has a high degree of interest in safety issues and, therefore, a motivation to report. It implies a degree of altruism in voluntarily passing on experience without necessarily expecting any direct improvement within the respondents' particular working environment.

5.3.3 Eliciting Information from Respondents

The previous section focused on the information that must be provided in order to elicit incident reports. In contrast, this section identifies information that forms must elicit from its participants.

Detection Factors

A key concern in any incident reporting system is to determine how any adverse event was detected. There are a number of motivations behind this. Firstly, similar incidents might be far more frequent than first thought but they might not have been detected. Secondly, similar incidents might have much more serious consequences if they were not detected and mitigated in the manner described in the report.

As before, there are considerable differences in the approaches adopted by different schemes. CIRS provides an itemised list of clinical detection factors. These include direct clinical observation, laboratory values, airway pressure alarm and so on. From this the respondent can identify the first and second options that gave them the best indication of a potential adverse event. It is surprising that this list focuses exclusively on technical factors. The web-based form enables respondents to indicate how teams help to resolve anomalies, however, it does not consider how the users' workgroup might help in the detection of an incident.

The local scheme of Figure 5.5 simply asks for the "Grade of staff discovering the incident". Even though it explicitly asks for factors contributing and mitigating the incident, it does not explicitly request detection factors. In contrast, ASRS forms reflect several different approaches to the elicitation of detection factors. For instance, the forms for reporting maintenance failures includes a section entitled "When was problem detected?". Respondents must choose from routine inspection, in-flight, taxi, while aircraft was in service at the gate, pre-flight or other. There is, in contrast, no comparable section on the form for Air Traffic Control incidents. This in part reflects the point that previous questions on the Air Traffic Control form can be used to identify the control position of the person submitting the form. This information supports at least initial inferences about the phase of flight during which an incident was detected. It does not, however, provide explicit information about what people and systems were used to detect the anomaly. Fortunately, all of the ASRS report forms prompt respondents to consider "How it was discovered?" in a footnote to the free-form event description on the second page of the report. In the ASRS system, analysts must extract information about common detection factors from the free-text descriptions provided by users of the system.

Causal factors

It seems obvious that any incident reporting form must ask respondents about the causal factors that led to an anomaly. As with detection factors, the ASRS exploits several different techniques to elicit causal factors depending on whether the respondent is reporting an Air Traffic Incident, a Cabin Crew incident etc. For example, only the Cabin Crew forms ask whether a passenger was directly involved in the incident. It is interesting that the form does not distinguish between whether the passenger was a causal factor or suffered some consequence of the incident. In contrast, the maintenance forms ask the respondent to indicate when the problem was detected; during routine inspection; in-flight, taxi; while aircraft was in service at gate; pre-flight or other.

In spite of these differences, there are several common features across different categories in the ASRS. For instance, both Maintenance and Air Traffic reporting forms explicitly ask respondents to indicate whether they were receiving or giving instruction at the time of the incident. Overall, it is surprising how few explicit questions are asked about the causal factors behind an incident. The same footnote that directs people to provide detection information also requests details about "how the problem arose" and "contributing factors". This is an interesting distinction because it suggests an implicit categorisation of causes. A primary root cause for "how the problem arose" is being distinguished from other "contributing factors". This distinction is not followed in the local scheme of Figure 5.5. The respondent is simply asked to identify "what contributed to the incident". The same form asks specifically for forms of equipment failure but does not ask directly about any organisational or managerial causes.

The web-based CIRS has arguably the most elaborate approach to eliciting the causes of an incident. In addition to a free-text description of the incident, it also requests "circumstantial information" that reveals a concern to widen the scope of any causal analysis. For instance, they include seven Likert scales to elicit information about "Circumstances: team factors, communication". Respondents are asked to indicate whether there was no briefing (1) up to a pertinent and thorough briefing (5). They must also indicate whether there was a major communication/co-ordination breakdown (1) or efficient communication/co-ordination in the surgical team (5). Again, such questions reveal a great deal about the intended respondents and about the people drafting the form. In the former case it reveals that the respondents must be aware of the general problems arising from team communications and co-ordination in order for them to assess its success or failure. In the latter case, such causal questions reveal that the designers are aware of recent literature on the wider causes of human error beyond "individual failure".

Consequences

Previous paragraphs have shown different reporting systems exploit different definitions of what constitutes an incident. These differences have an important knock-on effect in terms of the likely consequences that will be reported to the system. For instance, the distinction between the incident and accident reporting procedures of the FAA will ensure that no fatalities will be reported to the ASRS . Conversely, the broader scope of the CIRS definition ensures that this scheme will capture incidents that do contribute to fatalities. This is explicitly acknowledged in the rating system that CIRS encourages respondents to use when assessing the outcomes of an incident: Transient abnormality - unaware for the patient; Transient abnormality with full recovery; Potential permanent but not disabling damage; Potential permanent disabling damage; Death [464]. This contrasts with the local system that simply provides a free text area for the respondent to provide information about "what happened to the patient?". The domain dependent nature of outcome classification is further illustrated by maintenance procedures in the ASRS. Here the respondent is asked to select from: flight delay; flight cancellation; gate return; in-flight shut-down; aircraft damage; rework; improper service; air turn back or other.

The distinction between immediate and long-term outcomes is an important issue for all incident-reporting schemes. The individuals who witness an incident may only be able to provide information about the immediate aftermath of an adverse event. However, human 'error' and system 'failure' can have far more prolonged consequences. This is acknowledged in the Lack scale of prognosis used in the CIRS system. Transient abnormalities are clearly distinguished from permanently disabling incidents. The other schemes do not encourage their respondents to consider these longer-term effects so explicitly. In part this can be explained by the domain specific nature of consequence assessments. The flight engineer may only be able to assess the impact of an incident to the next flight. Even if this is the case, it is often necessary for those administering the schemes to provide information about long-term effects to those contributing reports. This forms part of the feedback process that warns people about the potential long-term consequences of future incidents.

Mitigating factors

Several authors argue that more attention needs to be paid to the factors that reduce or avoid the negative consequences of an incident [841]. These factors are not explicitly considered by most reporting systems. There is an understandable focus on avoiding the precursors to an incident rather than mitigating its potential consequences. For instance, the ASRS forms simply ask respondents to consider "Corrective Actions" as a footnote to the free text area of the form shown in Figure 5.7.

Similarly, the local form shown in Figure 5.5 asks respondents to describe “what factors minimised the incident”.

The CIRS again takes a different approach to the other forms. Rather than asking the user to describe mitigating factors in the form of free-text descriptions, this system provides a number of explicit prompts. It asks the respondent to indicate whether personal factors such as “appropriate knowledge, skill, experience or situational awareness” were recovery factors. The form also asks for information about ways in which equipment provision and team co-ordination helped to mitigate the consequences of the failure. Questions are also asked about the role of management and the working environment in recovery actions. Such detailed questions can dissuade people from investing the amount of time that is necessary to complete the 20 fields that are devoted to mitigating factors alone. Of course, the trade-off is that the other schemes may fail to elicit critical information about the ways in which managerial and team factors helped to mitigate the consequences of an incident.

Prevention

Individuals who directly witness an incident can provide valuable information about how future adverse events might be avoided. However, such individuals may have biased views that are influenced by remorse, guilt or culpability. Subjective recommendations can also be biased by the individual’s interpretation of the performance of their colleagues, their management or of particular technical subsystems. Even if these factors did not obscure their judgement, they may simply have been unaware of critical information about the causes of an incident. In spite of these caveats, many incident reporting forms do ask individuals to comment on ways in which an adverse event might have been avoided. The local system in Figure 5.5 asks respondents to suggest “how might such incidents be avoided”. This open question is, in part, a consequence of the definition of an incident in this scheme which included occurrences “that might have led (if not discovered in time) or did lead, to an undesirable outcome”. This definition coupled with the request for prevention information shows that the local system plays a dual role both in improving safety ‘culture’ but also in supporting more general process improvement. This dual focus is mirrored in the CIRS form:

> “What would you suggest for prevention of this incident? (check all appropriate fields): additional monitoring/equipment; improved monitoring/equipment; better maintenance of existing monitoring/equipment; improved arrangement of drugs; improved arrangement of monitoring/equipment; better training/ education; better working conditions; better organisation; better supervision; more personnel; better communication; more discipline with existing checklists; better quality assurance; development of algorithms / guidelines; abandonment of old ’routine’.”

This contrasts with the local system in which “complications which occur despite normal management are not critical incidents but if in doubt fill in a form”. Under the CIRS definition, failures in normal management would be included and so must be addressed by proposed changes.

The ASRS does not ask respondents to speculate on how an incident might have been avoided. There are several reasons for this. Some of them stem from the issues of subjectivity and bias, mentioned above. Others relate to the subsequent analytical stages that form part of many incident-reporting systems. An important difference between the ASRS and the other two schemes considered in this section is that it is confidential and not an anonymous system. This means that it is possible for ASRS personnel to contact individuals who supply a report to validate their account and to ask supplementary questions about prevention factors. CIRS does not provide direct input into regulatory actions. Instead, it aims to increase awareness about clinical incidents through the provision of a web based information source. It, therefore, protects that anonymity of respondents and only has a single opportunity to enquire about preventive measures. In the local system, the personnel who administer the system are very familiar with the context in which an incident occurs and so can directly assess proposed changes to working practices.

There has been a rapid growth in the use of incident reporting schemes as a primary means of preventing future accidents. However, the utility of these systems depends upon the forms that are used to elicit information about potential failures. This section, therefore, uses a comparative study of existing approaches to identify key decisions that must be made during the design of future documents. Much work remains to be done. At one level, the various approaches of the ASRS , CIRS and the local system have been validated by their success in attracting submissions. At another level, there is an urgent need for further work to be conducted into the validation of specific approaches. For instance, it is unclear whether techniques from the CIRS system might improve the effectiveness of the local system or vice versa. This work creates considerable ethical and methodological problems. Laboratory experiments cannot easily recreate the circumstances that lead to incident reports. Conversely, observation analysts may have to wait for very long periods before they can witness an incident within a real working environment. The lack of research in this area has led to a huge diversity of reporting forms across national boundaries and within different industries. We urgently need more information about the effects that different approaches to form design have upon the nature and number of incidents that are reported to these systems.

5.4 Summary

This chapter has argued that under-reporting continues to be a major limitation of most incident reporting systems. For instance, Barach and Small estimate that between 50 and 95% of medical incidents go unreported [66]. This problem is exacerbated by the difficulty of accurately assessing the nature and extent of under-reporting. For instance, most current estimates rely upon base-line estimates. These are derived by extrapolating the number of incidents that are observed within a narrowly defined sample set. Incidents can be identified by an exhaustive manual examination of the logs and records that are taken during a particular period of operation. Alternatively automatic and semi-automatic tools can be used to look for patterns in these data sets that might indicate a potential incident. However, both of these techniques are limited in that they cannot provide information about potential failures that were averted in good time. Nor can they provide information about many of the contextual and causal factors that are important when assessing the consequences of under-reporting. Observational studies avoid some of these problems but they tend to be expensive and controversial; workers may not agree to the independent monitoring of their daily activities. It is also difficult to identify the under-reporting of relatively low-frequency events using any of these techniques.

Subsequent sections went on to assess the strengths and weaknesses of mandatory reporting systems as a potential means of avoiding the problems of under-reporting. These systems, typically, enforce legal or administrative sanctions if individuals fail to report certain classes of incidents. However, recent clinical studies of reporting behaviour reveal that these systems are themselves biased towards high-criticality mandatory events or previously unseen adverse reactions. Mandatory systems are not, universally, effective in ensuring that contributors report more routine, low criticality incidents.

Automatic, real-time monitoring systems provide an alternative means of ensuring notification about adverse occurrences. It was argued that these tools often suffer from problems of precision or recall. Poor precision results in a high proportion of 'normal' incidents being identified as potential occurrences. These are often referred to as false positives. In contrast, poor recall occurs when many potential incidents go undetected by the system. However, it has also been argued that many of the barriers to these systems are not technical but social. For example, several groups have opposed the introduction of data logging equipment into the cabs of trains. It can also be difficult to interpret the causes of potential incidents that are detected by automated systems. It is for this reason that the NTSB and others have advocated the use of cameras to supplement flight data recorders. Again, however, there are strong and justified objections to what is partly seen as an intrusion on the rights of the crews who will be monitored.

Later sections of this chapter have examined a number of techniques that are intended to encourage greater participation within incident reporting systems. The decision whether or not to submit a report is affected by a number of considerations. In particular, the fear of retribution or disclosure may dissuade potential contributors. This fear can be addressed by trust in local champions or guarantors who both advocate and protect the system against external pressures. However, there is a danger that trust will be lost in the system if those champions are replaced. Contributions can also be encouraged if potential participants are reassured that their colleagues are contributing to the system. This is an important consideration if individuals fear that they may be perceived to be 'whistle blowers'. Similarly, it is important that potential participants know both what to report and how to report it. This is supported by various feedback mechanisms that provide examples of incidents that have already been investigated. These publications also reinforce the idea that contributions will be taken seriously and will be acted upon. Contributions can be encouraged by ensuring that the incident reporting system plays a visible part within wider management systems. Later chapters will stress the importance of integrating information about previous occurrences into both training practices and risk assessment procedures.

The closing sections of this chapter focused more narrowly upon the components of form design. A number of different approaches are considered. These include a paper-based local system that operates within a single hospital ward. They also include a national paper-based system that operates across the many different industries that support US aviation operations. These are, in turn, compared against an innovative Internet-based reporting system. The forms that are exploited by these schemes reflect different managerial and organisational constraints. For instance, the local scheme focuses on incidents that occur within the unit. It does not address managerial issues that cannot directly be influenced by staff within the unit. The national scheme does not face these limitation. Regulatory support ensure that structural issues can be addressed if they are raised as being significant by a number of different contributors. In contrast, the electronic system maintains the anonymity of each contributor and cannot, therefore, validate the information presented in each report. This places constraints on the sorts of follow-up actions that might be taken in response to each incident. It is concluded that more work is urgently needed to determine the detailed effects that such different strategies might have upon the success or failure of an incident reporting system. The lack of any objective data in this area is compounded by the lack of any published guidelines or advice on form design. As a result, there is a proliferation of local styles. Many of which needlessly repeat weaknesses that have been identified and corrected in the design of forms in other systems. For instance, many forms use terms such as 'slip', 'lapse' or even 'situation awareness' that continue to confuse potential contributors who have (sadly) never read the works of Reason or Rasmussen.

The following chapters explores techniques that can be used to investigate the causes of an incident once it has been detected. These include interview techniques that help investigators to take eye witness statements. They also include an outline table of contents for the preliminary reports that are used to inform others within an organisation in the immediate aftermath of a safety-related incident.

Chapter 6

Primary Response

The previous chapter looked at the problems that any incident reporting system faces in eliciting submissions about adverse occurrences or the potential for future accidents. The following sections build on this by looking at techniques that can be used to address the problems of gathering further information about an occurrence once it has been notified. These data gathering techniques produce the evidence that supports subsequent analysis. As a result they have an important impact on the outcome of any investigation. If necessary data is not secured then analysts may be forced to rely upon supposition and introspection. Similarly, if investigators obtain biased or partial information then the conclusions of an enquiry may not accurately reflect the underlying causes of an incident. Further problems arise because different approaches to data gathering obtain very different results. Later sections will examine the ways in which one-to-one interviews can provide very different accounts than peer group meetings. These potential problems are exacerbated by the difficulties of supporting an iterative approach to incident investigation. Often the subsequent analysis of an occurrence will help to identify the need for further information about the causes or mitigating factors that influenced an adverse occurrence. However, data may be lost, opinions and recollections may change over time, outside influences may affect the participation of key individuals. As a result, the answers that are obtained during subsequent investigations may not actually reflect the potential answers that might have been gathered during the initial stages of an enquiry.

Figure 6.1 again illustrates how these different generic phases contribute to the operation of an incident reporting system. This chapter, therefore, concentrates on phase B data gathering. This abstract model is intended to describe common features of many different reporting systems. The following quotation provides greater detail about the sorts of activities, listed as points 3 to 5, that contribute to data gathering in a medical incident report system. It also illustrates the way in which these activities depend upon the elicitation of reports, see points 1 and 2, and support the subsequent analysis of adverse occurrences, mentioned in points 5 and 6:

> "*Summary of investigation process*: All investigations consist of a series of steps that should be followed, as a matter of routine, when an incident is investigated:
>
> 1. Ascertain that a serious clinical incident has occurred and ensure it is reported formally. Alternatively identify an incident as being fruitful in terms of organisational learning;
> 2. Trigger the investigation procedure. Notify senior members of staff who have been trained to carry out investigations
> 3. Establish the circumstances as they initially appear and complete an initial summary, decide which part of the process of care requires investigation, prepare an outline chronology of events, and identify any obvious care management problems;
> 4. Structured interview of staff: Establish chronology of events; Revisit sequence of events and ask questions about each care management problem identified at the initial stage. Use framework to ask supplementary questions about reasons for each care management problem;

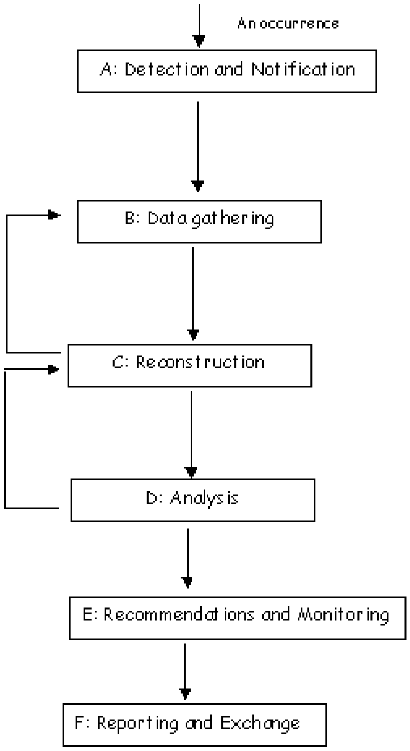
An occurrence
A: Detection and Notification
B: Data gathering
C: Reconstruction
D: Analysis
E: Recommendations and Monitoring
F: Reporting and Exchange

Figure 6.1: Generic Phases in Incident Reporting Systems

5. If new care management problems have emerged during interviews add them to initial list. Interview again if necessary
6. Collate interviews and assemble composite analysis under each care management problem identified. Identify both specific and, where appropriate, general contributory factors;
7. Compile report of events, listing causes of care management problems and recommendations to prevent recurrence
8. Submit report to senior clinicians and management according to local arrangements
9. Implement actions arising from report and monitor progress." [848]

It is important to emphasise, however, that individual reporting systems may different significantly from the blue-print provided by this list of activities. In particular, the opportunities for gathering further information are constrained by the procedures and practices that govern the management of any reporting system. The following paragraphs summarise the financial, social and technical issues that constrain data gathering exercises.

It may not be possible to identify the individuals who were involved in an incident. As a result, any subsequent data gathering must be based around teams or groups of individuals who might be involved in similar occurrences. Instead of interviewing the controller who was involved in a particular air separation violation, investigators must find other individuals who are willing to talk about the circumstances of previous incidents.

In a confidential system, it is likely that investigators will be able to identify the individuals or groups who reported an occurrence. However, this information may only be available during the initial stages of an enquiry. For instance, the UK CIRAS rail incident reporting system protects the identity of individuals by destroying all identifying information once an initial interview procedure has been completed. During those stages in which it is possible to identify the individuals who contributed a report, it is important not to compromise the confidentiality of the system. For example, requests to interview an operator can raise suspicions about the purpose of any enquiry. As a result, many confidential systems make contact with contributors outside of normal working hours. It should also be noted that such procedures place important restrictions on the gathering of confirmatory evidence. For example, it is difficult to interview the colleagues of a contributor without telling them the purpose of the meeting;

The architecture of an incident reporting system can also limit the opportunity for data gathering activities. For instance, the simple monitoring architecture described in Chapter 4 does not assume that there will be any further investigation of a particular occurrence. A report is received, an initial assessment is made about its relevance and then feedback about the incident is published. Such an approach is both simple to manage and cheap to operate. It can also reduce concerns about anonymity because no investigation is initiated. However, there are also important concerns about the reliability and completeness of the information that is contributed about each incident. The Swiss Internet-based CIRS system is an example of this architecture [756]. CIRS gathers information about occurrences in anaesthesia. It addresses many of the concerns, mentioned above, by exploiting a complex and detailed form that is intended to elicit as much information as possible when an occurrence is notified to the system. This approach relies upon the intellectual capabilities as well as the enthusiasm and commitment of potential contributors.

Chapter 2 introduced Leape's analysis of the comparative costs of incident reporting in different industries [480]. The Aviation Safety Reporting System spends about $ 3 million annually to analyse approximately 30,000 reports. This equates to about $100 (£66) per case. If this figure were applied to the 850,000 adverse events that are estimated to occur annually in the UK National Health Service, the cost of investigation would be £50 million per year. This would impose a considerable burden upon the service. Such burdens can most easily be considered in terms of the opportunity cost; do the benefits of this expenditure outweight the benefits of alternative investments that might have been made with this money?

Data gathering can also be limited by the availability of skilled personnel. As we shall see, interviewing personnel in the aftermath of an incident can be a non-trivial exercise. It is difficult to probe behind the filters of guilt or resentment that may colour an individual's response in the aftermath of an adverse occurrence. Similarly, the extraction of necessary technical information from automated logging equipment typically requires considerable expertise. The burdens imposed by these requirements are exacerbated when investigators must be drawn from a more limited pool of potential personnel. For instance, if a reporting system relies upon independent external organisations to conduct any initial data gathering then that agency may not have the necessary capacity to cope with any expansion in the scope of a system or with any changes in the level of participation.

As mentioned above, a high degree of technical skill can be required to extract and safeguard information from automated logging equipment. It should also be noted that technical limitations, including the granularity of information that can be recorded, also affect the results of any data gathering exercise. The recovery of technical data can also be compromised by management failures in the aftermath of an incident. For example, the flight data recorders (or 'black boxes') that are used to record flight parameters have relied upon loops of tape. In several incidents, these recorders have not been switched off after landing so that they have continued to record 'null' data over critical information about the course of an incident.

The remainder of this chapter looks at techniques that support data gathering in the aftermath of an incident within the limitations identified above. The analysis initially looks at the immediate response to an incident, including the requirement to safeguard the system. Later sections look at how investigators identify and acquire the information that supports the subsequent reconstruction and analysis of safety-related incidents.

As we have seen, there are many different ways in which an occurrence can be reported. For example, the staff who are involved in an incident might directly inform their managers that an adverse occurrence has taken place. Alternatively, an automated monitoring system might generate an alarm which, in turn, can initiate further data gathering. Information about an incident can also be provided by members of the public who may also have witnessed a potential failure. It is important that the managers of an incident reporting system should consider, and ideally support, these different possibilities if potential sources of notification are not to be ignored. In the following discussion, we will use the term ‘primary recipient’ to indicate the supervisors, managers or other nominated personnel who first receive an incident report. For instance, the UK Medical Devices Agency (MDA) requires that “local liaison officers” are appointed to perform this role [535]. In European Air Traffic Control, the primary recipient is typically the line manager or the supervisor of the officer who submits the report [423]. However, the primary recipient need not be employed by the same organisation as the contributor. In particular, they can be employed by an independent reporting agency, by the regulator or by some trade organisation. For instance, CIRAS staff are the first to receive notification of an incident from personnel who are employed by many different rail operating companies [197]. The term ‘primary recipient’, therefore, simply provides a place holder for the wide range of mechanisms that implement the duties which are described in this section.

Members of staff must understand the procedures that are associated with the immediate notification of an incident. For example, they must know how to pass information from the general public, from automated detection equipment or from their own experiences to the primary recipient. Such notifications are critical for occurrence registration. They warn primary recipients that report forms are being generated and that further data gathering may be required. Any delays in making this notification can jeopardise the acquisition of necessary information in the aftermath of an incident. There are also safety consequences if other systems are vulnerable to similar failures before any immediate remedial actions can be taken. Primary recipients must, in turn, warn others within their organisation. For example, they may be expected to inform higher levels of safety management. Many executives are embarrassed to learn of serious incidents from media enquiries rather than from the effective communication of safety concerns within their own organisation. In open reporting systems, it can also be good practice for primary recipients to brief other workers that an incident has taken place. Such actions are extremely important to preserve confidence in the reporting system; teams can see that some action is being taken. They can also elicit peer support for individual operators in the aftermath of an incident. Finally, it is often important to warn other organisations with a ‘stake’ in any incident investigation. For instance, air traffic control reporting procedures often contain a list of contacts and telephone numbers that should be called in response to particular occurrences. For example, if an incident involves a military flight then information should be passed to the force’s duty liaison officer. If an incident has implications for other sectors operated by other national organisations then they also might be alerted to a potential investigation.

It is possible to envisage a number of circumstances in which personnel might not want to submit occurrence reports to the groups and individuals who are normally nominated as ‘primary recipients’. For example, there is an understandable reluctance to provide reports that might jeopardise an individual’s relationship with their immediate superiors, especially if those superiors are implicated by an occurrence. Special provision should be made for such circumstances. However, previous comments about anonymity and the problems of under-reporting indicate that such channels may not be used very frequently unless the supervisor or manager’s behaviour has become irredeemable. The difficulties faced by junior personnel in questioning and reporting the ‘errors’ of their seniors can be illustrated by incidents drawn from the aviation industry. Crew Resource Management (CRM) training has been introduced to explicitly help staff overcome their inhibitions in intervening to question the actions of their seniors. The following incident illustrates how this training can fail to have a sufficient impact on operator behaviour:

> “*(Editorial comment) Recognition of the potentially hazardous effects (of the flight deck gradient) is often included as an aspect of CRM training, but the problem can be extremely complex, particularly if combined with an apparent short-term incapacitation. In such circumstances, it is often difficult for the junior crew member to intercede.*
>
> It was the Captain’s leg. He is an experienced pilot, capable and well liked and in no way overbearing. On short finals to Runway 30 at ####, after a good, stabilised visual circuit and approach, the aircraft begins to descend below the Visual Approach Slope Indicator (VASI) indications, giving finally four reds. As the runway has a displaced threshold and the obstacle was now behind us I make no comment, as I presume the descent (below the correct glide-path) is intentional to facilitate an early touch-down point. The Captain now sees the VASI indications, says so, and applies power. I call

> 'Rad Alt 50', '30' and '20' but we don't land. I inform the Captain we are floating and to put the aircraft on the ground. He seems surprised by my call, but removed power and lands. However, we are between a third and a half of the way down the runway. The Captain appears transfixed by the runway and hasn't engaged reversers as per SOP. I call for reversers and query the autobrake setting of level three out of five available levels. He makes no response although he is not obviously unwell. I state that I am increasing autobrake to level four. He doesn't acknowledge. As speed reduces he finally deploys the reversers, but as our Normal Operations Standard Operating Procedures, only at idle thrust. We stop with approximately 200 feet runway remaining. On taxi back he states he had difficulty reading the VASI and no other discussion occurs. With hindsight I allowed my attitude of respect and friendliness toward the Captain to influence my actions. I was insufficiently assertive once the incident was in progress and prior to the incident I presumed rather than checked the reasons for his flight profile." [173]

This incident report illustrates how individuals still fail to question the actions of their colleagues even when they believe that their safety and the safety of their passengers might be threatened. This failure is all the more remarkable given that CRM training deliberately includes help in recognising when to question such behaviours. Given such reluctance it should not be surprising that very few contributors will use alternative procedures to implicate the normal 'primary recipients' of incident reports. The previous quotation is, however, more complex than this analysis suggests. It illustrates the way in which individuals may contribute information about an incident even though they failed to question their colleagues actions during the occurrence itself. It can, therefore, be argued that it illustrates the importance of providing alternative reporting mechanisms. If these procedures are used then it provides mixed news about the wider safety of any system. On the one hand, it may indicate a strong safety culture in which individuals are happy to question the actions of their colleagues. On the other hand, these reports are disappointing if subsequent analysis indicates potential problems with the behaviour of 'primary recipients' who play an important role in the success of any incident reporting system.

6.1 Safeguarding the System

Figure 6.2 illustrates part of the checklist that is to be used whenever US Army commanders receive notification of an accident or incident [806]. This check-list is intended to support their actions in the crucial first hours after an adverse occurrence has been reported. This initial response is critical because the primary recipient must act both to safeguard their system and to protect any necessary data about the course of an incident. As can be seen, the first items on the checklist are to secure the site of any incident and to obtain witness statements. Further items cover survey procedures and the notification of relevant authorities that an incident or accident has occurred.

The primary recipient's first responsibility is to safeguard any systems that are involved in an incident. Chapter 4 has described how the operators who are involved in an incident are often removed from further operation. This provides them with an opportunity to gather their thoughts, to document the events leading to the incident and to complete an initial report form. It also removes the additional fear of committing further 'errors' in the aftermath of an incident. Regulators often view such compound failures as indicative of a failure by the primary recipients to take adequate measures to safeguard the system.

6.1.1 First, Do No Harm

It is critical that any remedial actions should not exacerbate the consequences of any initial failure. This is, however, a non-trivial requirement. Feelings of guilt or loyalty can encourage individuals to take ill-advised risks in the aftermath of an adverse occurrence. Inadequate training, incomplete information about the nature of an incident or the potential impact of their actions can all predispose 'primary recipients' to act without adequate forethought. The following quotation illustrates many of these issues. It is important to note that the investigators were anxious both to praise the initiative

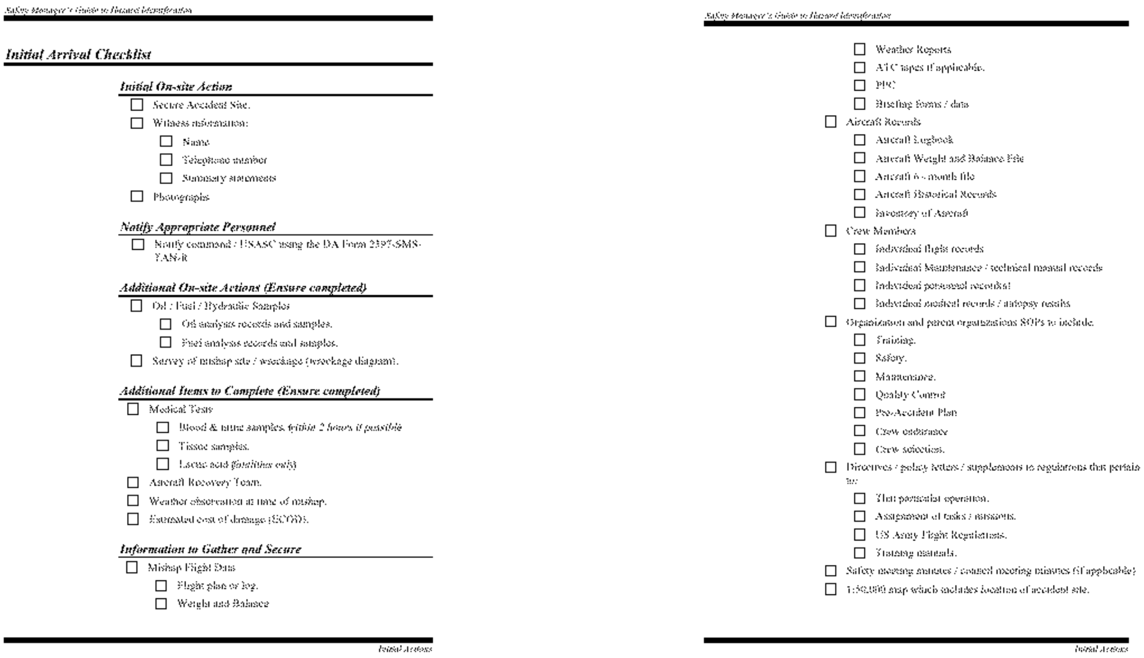

Safety Manager's Guide to Hazard Identification

Initial Arrival Checklist

Initial On-site Action

- □ Secure Accident Site.
- □ Witness information:
  - □ Name
  - □ Telephone number
  - □ Summary statements
- □ Photographs

Notify Appropriate Personnel

- □ Notify command / USASC using the DA Form 2397-SMS-YAN-R

Additional On-site Actions (Ensure completed)

- □ Oil / Fuel / Hydraulic Samples
  - □ Oil analysis records and samples.
  - □ Fuel analysis records and samples.
- □ Survey of mishap site / wreckage (wreckage diagram).

Additional Items to Complete (Ensure completed)

- □ Medical Tests
  - □ Blood & urine samples *(within 2 hours if possible*
  - □ Tissue samples.
  - □ Lactic acid *(fatalities only)*
- □ Aircraft Recovery Team.
- □ Weather observation at time of mishap.
- □ Estimated cost of damage (ECOD).

Information to Gather and Secure

- □ Mishap Flight Data
  - □ Flight plan or log.
  - □ Weight and Balance

Initial Actions

Safety Manager's Guide to Hazard Identification

  - □ Weather Reports
  - □ ATC tapes if applicable.
  - □ PPC
  - □ Briefing forms / data
- □ Aircraft Records
  - □ Aircraft Logbook
  - □ Aircraft Weight and Balance File
  - □ Aircraft 6 - month file
  - □ Aircraft Historical Records
  - □ Inventory of Aircraft
- □ Crew Members
  - □ Individual flight records
  - □ Individual Maintenance / technical manual records
  - □ Individual personnel records(s)
  - □ Individual medical records / autopsy results
- □ Organization and parent organizations SOPs to include.
  - □ Training.
  - □ Safety.
  - □ Maintenance.
  - □ Quality Control
  - □ Pre-Accident Plan
  - □ Crew endurance
  - □ Crew selection.
- □ Directives / policy letters / supplements to regulations that pertain to:
  - □ That particular operation.
  - □ Assignment of tasks / missions.
  - □ US Army Flight Regulations.
  - □ Training manuals.
- □ Safety meeting minutes / council meeting minutes (if applicable)
- □ 1:50,000 map which includes location of accident site.

Initial Actions

Figure 6.2: US Army Preliminary Incident/Accident Checklist

and endeavour of the crew but also to point out the potential consequences of an ill-considered response to an adverse occurrence:

> The 2nd Mates muster station was on the first bridge deck, the helicopter landing area, and his fire duty to take charge of a fire team. However, after the evacuation of the engine room, the 2 nd Mate took it upon himself to go alone and search for the 3 rd Engineer. He mistakenly understood that the 3 rd Engineer was still making his way out of the shaft tunnel. Although he advised the bridge by radio of his intended actions, he had no breathing apparatus and nobody was standing by to assist him. He went alone down the vertical after tunnel escape, along about 20 m of the shaft tunnel and into the engine room. He did not know if the atmosphere in the shaft tunnel was safe and, more particularly, whether the atmosphere in the engine room could support life. Fires deplete oxygen and, although the fire would have drawn air through the shaft tunnel, the combustion of fuel and the breakdown of insulation produces poisonous gases. The Inspector acknowledges that the 2 nd Mates actions were well-intentioned but he could easily have fallen, become disorientated or overcome by smoke, thereby hazarding the lives of any search party and compromising the fire fighting effort...
>
> *Conclusions.* In general the response to the fire by the ships crew and the expeditioners on board was measured, effective, demonstrated initiative and reflects great credit to all on board. Entry into any area adjacent to a fire, however, alone and without breathing apparatus or backup, is extremely hazardous and could compromise an entire fire-fighting effort. [48]

'Primary recipients' must address a number of complex problems in order to safeguard complex application processes. These problems are determined both by the nature of the failure and by the support that is afforded by remaining protection systems. For example, automatic deluge systems can quickly establish control over a reported fire. Similarly, critical tasks can be delegated to other members of a crew if an incident indicates excessive workload for key individuals. A drug mis-administration error may require both immediate and long term intervention to stabilise the patient's condition. It is important to remember, however, that any subsequent intervention must not exacerbate an adverse occurrence. A number of factors help to determine whether such reactions are likely to safeguard the continued operation of a complex system:

- *poor training.* Many industries have drafted guidelines that are intended to ensure that personnel are trained in emergency response techniques. Many of these guidelines focus on the need to ensure that skills are reinforced through simulated exercises. There are many potential problems, however. It can be difficult to organise simulations that involve representative of the different groups that must coordinate their activities in the aftermath of an incident [745]. There are considerable barriers to such joint simulations. These include organisational and financial constraints. They also include the underlying problems of ensuring a common 'mental model' both of the nature of any potential incident and the best means of addressing it [217]. There is also a danger that simulations may not reflect the challenges posed in the aftermath of an actual incident [874]. One means of ensuring that simulations do reflect potential failures is to ensure that they are based upon accurate observations of the previous failures that have been submitted to incident reporting systems.

- *situation awareness.* Chapter 3 described the general problems that arise when individuals and teams must continually predict and respond to changes in application processes. Interruptions, high-workload and a myriad of other 'performance shaping factors' jeopardise accurate assessments of the current and future states of complex systems. This creates particular problems if individuals must respond to incidents that resulted from a loss of situation awareness. If the primary recipient has been called from other duties then they must quickly assess the state of the system. However, any information that they gain from the operators will reflect their initial loss of situation awareness. It is likely to be incomplete and possibly inconsistent. This can have an adverse effect on any subsequent intervention by the primary recipient.

- *time pressure.* Time pressures compound the problems of accurately assessing the state of a system prior to any response to a reported occurrence. As with many aspects of incident reporting, the precise nature of these pressures will vary from domain to domain [437]. In air traffic management, air proximity warnings must be resolved almost immediately if collisions are to be prevented. In other domains, such as batch chemical processing, operators may have minutes and even hours to rectify an adverse occurrence. There are two different dangers associated with time pressures in the immediate response to an incident and both are closely related to the more general problems of situation awareness, mentioned above. Firstly, if a process changes gradually over time then it may be difficult for people to notice slowly developing trends that emerge over many hours [438]. Secondly, in processes that require rapid intervention there is a danger that personnel will intervene before they understand the true nature of the problem at hand. Several regulatory agencies have responded to these different pressures by requiring that operator wait for some specified period of time, or that they ensure agreement with their colleagues, before actively intervening in the aftermath of an incident. Duncan describes how such measures have created delays that, in turn, have threatened the safety of a number of nuclear systems [219].

- *lack of information.* In order to act effectively to safeguard any system, it is important that the primary recipient of any incident report can rapidly access relevant information. This includes details about the state of the system prior to the incident and information about any interaction with an application as the incident develops. It also includes accounts of any initial actions that staff may have taken to mitigate the immediate effects of an adverse occurrence. This is particularly important in the medical domain when the patient's reaction to these interventions provides important guidance for further remedial actions. In consequence, both the medical and aviation industries specify protocols and procedures that govern the passing of information following particular incidents. When these protocols are broken then these is a considerable danger that the primary recipients will fail to recognise the nature of the incidents that they must address [10]. However, it is important not simply to consider ways in which this information can be made accessible to primary recipients. It is also critical that they are trained to avoid problems of interpretation and analysis, such as the confirmation bias that can impair an individual's ability to consider alternative hypotheses. Later sections in the Chapter will consider the sister problems of recognition and judgement bias that can also impair the primary recipient's ability to use information in the aftermath of an incident.

- *lack of system support.* The primary recipient's ability to safeguard their system is, at least partially, determined by the level of available system support. An incident can often compromise their ability to intervene effectively. In many situations this forced them to resort to ad hoc measures or deliberate fall-back mechanisms to retrieve the situation, in other incidents they are not so fortunate:

  > "During the flight, the en route air traffic controller inadvertently cleared the aircraft to descend to an altitude that was below the minimum vectoring altitude (MVA) for the area. The MVA is the lowest altitude that meets obstruction clearance requirements in the specified airspace, and is the lowest altitude that Transport Canada has approved for vectoring of aircraft by air traffic control (ATC). The crew of ABL814 accepted the clearance and descended. By the time the controller recognised the problem, the aircraft had descended below radio coverage and could not be contacted directly using NAV CANADA's ground-based communications network". [620].

  This incident illustrates how those who are involved in an incident can use alternate safety systems to mitigate the consequences of an initial failure. However, the same system limitations affect the primary recipients who must also use the available infrastructure to safeguard their system.

- *need to preserve levels of service.* The primary recipient's ability to intervene to safeguard their system can also be constrained by external pressures to maintain particular levels of service. This raises particular problems, as we shall see in later sections, when primary recipients must both protect evidence of a failure and yet also enable the system to continue to operate. This is illustrates by the UK MDA's regulations for incident reporting:

  > "Defective items should not be repaired (either in-house or by a third party), returned to the manufacturer/supplier or discarded before an investigation has been carried out. The manufacturer or supplier should be informed promptly, and allowed to inspect the items if accompanied by an appropriate person... If devices are required to be kept in use, where possible remove defective parts so that the equipment may be repaired for re-use. Any parts so removed must be quarantined and securely stored pending investigation. MDA's advice should be sought and, in all cases, the defective parts should be clearly identified and kept secure. If it is not possible to remove defective parts or withdraw the machine from use, staff should be made aware of the need for increased vigilance and extra caution during use (see Evidence below). [535]

  At first site, it might seem that remedial actions must take priority over such 'quality of service' issues. However, the denial of Air Traffic Management services is likely to create further incidents and accidents. In other situations, poor situation awareness, the lack of necessary information and inadequate system support can place key individuals in an invidious position. For example, Offshore Installation Managers initially decided not to shut-down production on connected installations following the initial reports of fire on the Piper Alpha [193]. This had significant consequences because these inter-connections enabled gas to continue to escape from ruptured pipes on the Piper Alpha. If they had shut down production it would have caused an "almost immediate reduction in the flow of oil that was fuelling the fire in the centre of the platform".Their decision was justified because they had reason to believe that the Piper's on-board systems could cope with the emergency.It took a number of communications with company representatives and their fellow installation managers before the decision was taken to shut-down production. Their response was delayed not simply by a desire to continue production during what they believed to be a controllable incident, it was also exacerbated by the failure of communications systems. This incident illustrates how several of the factors in this list can combine to delay or frustrate an effective response to adverse occurrences. Emergency planning and disaster management programmes, mentioned above, are specifically designed to help staff cope with the 'wicked' problems posed by such compound failures.

6.1.2 Incident and Emergency Management

The previous list mentioned that many organisations compile detailed plans for incident management. These are then rehearsed during simulated rehearsals. The amount of guidance that is provided for the compilation of these plans varies from industry to industry. The level of guidance also varies between particular types of incidents within the same industry! For example, these are very few national standards that guide the immediate response to iatrogenic injuries. In contrast, the UK MDA [535] and the US Food and Drug Administration (FDA) [272] issue detailed guidance on the primary recipients duties in response to reports of equipment failures. The degree to which emergency procedures are integrated into wider safety management practices also varies considerably. For instance, the following guidelines present the International Maritime Organisation's requirements for the integration of contingency planning into shipboard safety management systems:

> "The Guidelines provide a framework for preparing an emergency response plan to deal with emergency situations. The International Safety Management code requires contingency planning as part of the ship's Safety Management System (SMS). The Guidelines set out a modular designed structure for contingency planning which provides a quickly visible and logically sequenced source of information and priorities which can reduce error and oversight during emergency situations. The system should be applied to each individual ship, taking into account ship type, construction, cargo, equipment, staffing and route. A typical system would include six modules:
>
> Module I: Introduction - providing guidance and an overview;
>
> Module II: Provisions - should contain information and explanations for the development of the system based on the suggestions for improvement gained from the individual company and shipboard personnel;
>
> Module III: Planning, - preparedness and training should provide for emergency training and education of shipboard personnel to develop general awareness and understanding of actions to be taken in the event of an emergency;
>
> Module IV: Response actions - should provide for emergency training and education of shipboard personnel to develop general awareness and understanding of actions to be taken in the event of an emergency, including potential emergency situations;
>
> Module V: Reporting procedures - the System must specify procedures for making the initial report to the parties concerned since any ship involved in an emergency situation, or in a marine pollution incident, will have to communicate with the appropriate ship interest contacts and coastal State or port contacts;
>
> Module VI: Annex(es) - other requirements." [389]

Such general requirements can be supplemented by special provisions that guide intervention in the aftermath of particular types of incident. In other words, the development of an incident response plan must be guided by risk assessment techniques. Clearly, more detailed provising ought to be made for higher risk incidents. For example, the IMO issues special regulations to govern emergency procedures for ships carrying irradiated nuclear fuel (INF) [386]. Ships transporting these materials must develop shipboard emergency plans that include the procedure to be followed in reporting an incident involving INF materials. They must also have prepared a list of the authorities to be contacted in the event of an incident. They must have compiled a checklist of action to be taken immediately to "prevent, reduce or control the release of INF Code materials". Finally, contingency plans must describe procedures and points of contact for co-ordinating with local and national authorities. Such general requirements can also be supplemented with more detailed guidelines about the sorts of incidents that should be explicitly considered within a contingency plan. Problems arise, however, from the difficulty of predicting the precise types of incidents that will be arise. Later sections will go on to argue that it is often difficult for primary recipients and their colleagues to accurately assess the potential risk of an incident in its immediate aftermath. For now it is sufficient to realise that the level of detail required in a contingency plan, in part, reflects the degree of risk associated with the consequences both of the potential incidents and of a failure to adequately deal with those incidents.

The previous paragraphs have described how the primary recipient must safeguard the system following an incident report. The problems of gathering information and of assessing the severity of an incident combine to make it likely that such responses will be error prone and may even exacerbate an adverse occurrence. As a result, many organisations codify procedures for the initial response to an incident in the form of emergency management systems. There are, however, a number of additional factors that complicate attempts to safeguard application processes in the aftermath of an incident. For example, some regulatory bodies use the immediate response to an incident as one means of measuring its criticality. This raises a number of complications, for example when the response to an incident is based upon a precautionary approach in which the primary recipient ensures the safety of the system by assuming the 'worst case' scenario. This has led the US Federal Railroad Administration to explicitly state the extent to which precautionary treatment can be taken into account when assessing the severity of an adverse occurrence:

> "Treatment provided in response to an event such as a dog bite may be precautionary. For example, a rabies shot following a dog bite is precautionary treatment, so the injury would be reportable. The single stated exclusion to reporting injuries which require precautionary treatment is a tetanus shot, since the decision to give this shot is generally based on the date of the last injection rather than the severity of the injury. Under certain circumstances some treatments occurring prior to a diagnosis may not, by themselves, make a case reportable. For example, it is often a standard procedure of emergency rescue teams to administer preventive treatment such as oxygen or apply an intravenous saline solution while a patient is being transported to a medical facility for further evaluation. Such preventive treatment does not make the injury reportable." [233]

It is important to emphasise that the primary recipient's actions in safeguarding the system are unlikely to provide adequate long-term fixes. Testing is required in order both to ensure that any remedial action actually does protect against the recurrence of an incident and that any recommended fixes do not introduce unwanted side-effects that may themselves threaten safe and successful operation. Typically, any longer term changes to the design or operation of a system must be documented and justified through changes in any supporting safety case that is approved by a regulator [434]. Further actions are also required if investigators are to determine whether particular fixes are adequate for similar systems in other plants or operating conditions. Further information about the causes of an incident often creates the need to implement additional remedial actions. In particular, the primary recipients view of a single incident must be placed in the context of any previous incidents with similar causes or consequences. These concerns make it likely that any initial actions in safeguarding the system are unlikely to provide long-term solutions:

> "The discovery that a remedial action is necessary may be a direct result of one or more medical device adverse event reporting (MDR) reportable events occurring, or may be discovered through the performance of internal analyses using appropriate statistical or other acceptable methodologies. Action taken to fix a single device involved in an MDR reportable event is not remedial action." [258]

In an ideal world, there would be a point in time when the primary recipient is confident that they have ensured the continued safety of their system. This would enable them to start acquiring additional logs and eye-witness statements about an adverse occurrence. In practice, however, data gathering activities are likely to be punctuated by the knock-on effects of their immediate actions in the aftermath of an incident. For example, high workload incidents often force managers to reallocate tasks to other members of staff. This creates the potential for further incidents until 'normal' working patterns are resumed. However, there is still the potential for further incidents to occur before long-term changes can be implemented. Similarly, back-up systems are typically less reliable than the primary systems that they replace [762]. In consequence, primary recipients can find themselves under a considerable amount of stress as they struggle to coordinate the initial response to an adverse occurrence.

6.2 Acquiring Evidence

The primary recipient of an incident report is, typically, responsible for ensuring that any relevant evidence is secured in the aftermath of an incident. This raises the problem of defining what is, and what is not, relevant to the course of any future investigation. The United States Federal Rules of Justice (Article II, Judicial Notice) define relevant evidence to mean "evidence having any tendency to make the existence of any fact that is of consequence to the determination of the action more probable or less probable than it would be without the evidence". In other words, evidence helps in the determination of fact.

6.2.1 Automated Logs and Physical Evidence

The importance of any fact cannot easily be predicted in the immediate aftermath of an incident. For example, flight data recorders are routinely inspected in the aftermath of an incident. ICAO requirements specify that effective use shall be made of flight recorders in the investigation of an incident (Annex 13, Section 5.8 [384]). However, this does not necessarily mean that this source of data will actually be useful in any subsequent investigation:

> "The flight recorders fitted to both aircraft were not removed for analysis. Adequate data for the investigation was available from the recordings of Air Traffic Control Radio Telephony frequencies and secondary radar returns." [14]

The difficulties in predicting precisely what evidence will be relevant to any investigation has led a number of organisations to publish check-lists that specify the sources of data that must be secured in the aftermath of an adverse occurrence. These documents must embody international agreements, such as ICAO Annex 13 mentioned above. They must also meet national and institutional guidelines that are intended to specify minimum standards across comparable organisations. However, it is important that such checklists also explicitly identify any local systems that might provide useful information about an incident. For example, the following sources of information must be gathered if an incident is reported to EUROCONTROL's Upper Air Control Centre in Maastricht: Recordings of system data (including PCPAMPLAY, PAMFLG, PAMTRK, PAMPOS); Voice Recordings; Statement by staff involved; DCFEP Recordings; Daily Log entries; Position Log, Break Lists and Shift Rosters; Personal Databank Information - ATC Related; Eurocontrol Operations Manual Part 1 and 2; Systems Manual Maastricht UAC; Internal Notes, Briefing Sheets and Attachments to Briefing Sheets; Supporting Technical Information; Letters of Agreement; National Documentation; ICAO Documentation [68]. As can be seen, the safety manager who compiled this took considerable care to enumerate the local systems that must be inspected to provide the data that is required by the ICAO and recommended by EUROCONTROL's Safety Regulation group.

Such lists can be deceptive. They hide the practical difficulties that primary recipients have to address in order to gather necessary data .

> "The first reported tampering with an event recorder was noted in the investigation of a 1982 side collision of two freight trains near Possum Grape, Arkansas. A deadheading conductor stated the speed-recording device was working properly prior to the accident; but several hours after the accident, a railroad official found the case broken open and the tape missing, even though the locomotive cab had not been damaged." [214]

Much has been done to improve the crash-worthiness of these recording devices and their logs. However, there may still be considerable personal danger involved in taking the necessary actions to safeguard automated logs.

> "In the 1994 investigation of rear-end collision of between a moving freight train with a standing freight train at Cajon, California, the Safety Board again found that 3 of the 4 solid state multi-event recorders had been destroyed by fire
> indexData recorders!limitations. Only the carriers quick action to remove the data pack, as the fire approached the locomotive, salvaged the fourth event recorder, which provided important data for the investigation. In June 1997 two freight trains collided and derailed in Devine, Texas. All of the event-recorder data were lost because impact forces or fire, or both destroyed the recorders. The Safety Board issued Safety Recommendation R-98-030 to the Federal Railroad Administration, asking them to develop and implement event recorder crashworthiness standards for all new or rebuilt locomotives by January 1, 2000." [214]

As we have seen, the necessity of safeguarding a system can also delay an efficient response. There can also be bureaucratic and technical barriers to data collection. It is important that these are minimised within any emergency response plan. If such issues are not addressed then these is a danger that necessary data can be destroyed, repaired or deleted. For example, many cockpit voice recorders (CVR) rely upon solid state storage devices that have enough capacity to hold approximately thirty minutes of conversation. Previous recordings are continually erased in order to make space for current data. As a result, if the recording is not halted in the aftermath of an incident then the CVR will be over-written. The report into the Puerto Plata air accident illustrates how a failure to safeguard critical data can occur even in the aftermath of major incidents. The need to motivate train staff to ensure the protection of necessary evidence is correspondingly greater for less critical incidents:

> "The CVR, which was of thirty minute recording duration, had been allowed to continue to operate after the aircraft had landed. This, together with the diversion flight from Puerto Plata, ensured that the audio recorded during the accident flight had been over-written. It thus proved to be of no use to the investigation. " [16].

This is one of a large number of similar incidents in which CVR data has been lost [19] . This incident is instructive for other reasons. In particular, it suggests that any attempts to introduce cockpit video monitoring, as described in Chapter 5, must also consider effective procedures for protecting such recordings once they have been made. Partly as a result of these concerns, the ICAO have initiated a campaign to increase the duration of CVR devices from thirty minutes to two hours [383]. Further problems complicate the primary recipient's task of collecting evidence about the causes and consequences of adverse occurrences. In particular, the increasing development of heterogeneous and distributed systems makes it highly likely that any data acquisition will depend upon the cooperation of several different organisations. In the immediate aftermath of an incident, the primary recipient may be able to do little more than alert their colleagues that some of their logs and transcripts must be saved. However, as time goes on they or other appointed investigators will have to collate the information from these disparate sources. For example, the European Turbulent Wake Incident Reporting System initially received forms from pilots that detailed the type of aircraft involved, its position, flight phase and control settings [547]. The pilot also assessed the effect of the vortex on the aircraft. They could submit a shortened version of the form if they could back up their submission with Flight Recorder information. After receiving notification of an incident, the primary recipients would obtain information about leading and following aircraft form the Air Traffic Service providers. This together with terminal radar data was used to verify the position and separation of the aircraft involved. Meteorological data was also collated in response to an incident report using the METAR reports that are made every half hour at all terminals during operating hours. The METARs immediately preceding and succeeding the incident provide information about wind, temperature, cloud cover, humidity and visibility. This brief description reveals that for every potential windshear incident the primary recipient would have to collate information from the pilot, from their data recorders, from en-route and terminal air traffic control systems and from meteorological records.

This section has reviewed some of the problems that arise when the primary recipient of an incident report, such as the "local liaison officer" for the MDA [535] or Air Traffic Management supervisor [423], must safeguard necessary evidence. These problems include the need to meet national and international requirements for the collection of data in the aftermath of particular incidents. In order to do this they must ensure that automated logs are not deleted or corrupted. They must also ensure the cooperation of their colleagues in other agencies who often control other sources of corroborative information. Previous paragraphs have, however, focussed on the collection of data from automated sources. There are many other potential sources of evidence that must be protected in the aftermath of an incident. These can have create some particular problems for the primary recipients. For example, in order to meet ICAO requirements they must ensure that accident and incident investigators have "unhampered access to the wreckage and unrestricted control over it to ensure that a detailed examination can be made without delay by authorised personnel participating in an investigation" [384]. Pragmatically this can force the primary recipients of an incident report to instigate police and crowd control measures to preserve the physical evidence associated with severe near-miss incidents. In the medical domain, the collection of physical evidence raises even more complex issues. For instance, contaminated equipment must be labelled and kept in some form of quarantine. If this is not possible, then the state of the device at the time of the incident must be recorded by any and all means available for that it can be reconstructed during an investigation [535]. The following excerpt illustrates these concerns and recommends means of ensuring that physical evidence is protected in the aftermath of an incident:

> *Contaminated items.* "Where decontamination/cleaning would destroy vital evidence, the item should be placed in protective containment, labeled and placed in quarantine. MDA and the manufacturer/supplier should be contacted for advice prior to any further action being taken. IT IS ILLEGAL TO SEND CONTAMINATED ITEMS THROUGH THE POST
>
> *Evidence.* All material evidence should be labeled and kept secure. This includes the products themselves and, where appropriate, packaging material or other means of batch identification. The evidence should not be interfered with in any way except for safety reasons or to prevent its loss. If necessary, a record should be made of all readings, settings and positions of switches, valves, dials, gauges and indicators, together with any photographic evidence and eye-witness reports. If it is believed that an urgent examination of the defective item (or related items) is needed, then consideration should be given to sending the item(s) to MDA's Adverse Incident Centre, or inviting MDA's device specialists to inspect them on site." [535]

This quotation illustrates the emphasis that many regulators place upon documented procedures for the handling of physical evidence and automated logs. This information is critical to the success or failure of any subsequent attempts to reconstruct an incident or analyse its causes. MDA requirements also include detailed instructions that restrict the primary recipients interaction with product manufacturers. They are entitled to provide them with samples of unused stock from a large batch of similar products. However, they must ensure that manufacturer are not be allowed to "exchange, interfere with, or remove any part of the product" implication in an incident if it could prejudice subsequent investigations [535]. Such concerns are not simply based upon a natural desire to support the causal analysis of any incident. Legal consideration affect the ways in which evidence is handled in the aftermath of an adverse occurrence. ICAO requirements explicitly consider some of the problems that this creates. For example, possible 'conflicts' between investigating and judicial authorities regarding the custody of flight recorders and their recordings "may be resolved by an official of the judicial authority carrying recordings to the place of readout, thus maintaining custody" [384]. Previous sections have mentioned that investigatory bodies, such as the UK Air Accidents Investigation Branch (AAIB), determine the circumstances and causes of accidents and incidents rather than apportion blame or liability. However, their findings are often used in subsequent litigation. Similarly in no-blame incident reporting systems, such as the ASRS , there is still the possibility that an incident report may trigger a criminal prosection that will depend upon the primary recipient's ability to safeguard necessary evidence. As a result, it is important that the techniques that are used in gathering and protecting evidence should be beyond reproach.

The primary recipient of an incident report must not simply collate and safeguard data for any subsequent investigation. They must also ensure that this data is protected from (ab)use by unauthorised individuals and organisations. The information that they acquire will be extremely sensitive for the people involved in the incident and for the organisations that they represent. This evidence can also have important implications both for regulatory authorities and, increasingly, for political administrations. Much of this sensitivity stems from public and media interest in incidents and accidents. As a result, many organisations argue that strong sanctions must be taken against individuals who 'leak' information before the publication of an official report. In addition to these more general concerns, there is a particular sensitivity about the release of data and voice recordings in the aftermath of aviation incidents. This stems from the ethical issues that are raised by attempts to broadcast the last actions of crews who are struggling to ensure the safety of their passengers [302]. Even in less serious incidents, there is a strong concern that the disclosure of evidence to the media or other partial sources could jeopardise confidentiality. Unless such disclosures are prevented then the natural fear of retribution will dissuade individuals from contributing to a system. As a result, international regulations have been drafted to explicitly restrict the disclosure of any information that is gathered by the primary recipient and other investigators in the aftermath of an adverse occurrence:

> " *5.12 Disclosure of Records* The state conducting the investigation of an accident or incident, wherever it occurred, shall not make the following records available for purposes other than accident or incident investigation unless the authority responsible for the administration of justice in the State determines that their disclosure outweighs the adverse domestic and international impact such action may have on that or any future investigation: all statements taken from persons by the investigation authorities in the course of their investigation; all communications between persons having been involved in the operation of the aircraft; medical or private information regarding persons involved in the accident or incident; cockpit voice recordings and transcripts from such recordings; and opinions expressed in the analysis of information including flight recorder information." [384]

Not only must primary recipients be aware of their duties of confidentiality, it can also be important for everyone involved in gathering evidence to understand how it may contribute to any subsequent legal proceedings. As mentioned previously, in the early stages of an investigation it may not be apparent whether an incident involves a criminal act. Even if the incident itself does not directly fall under the criminal law, evidence that is gathered in the aftermath of an adverse occurrence can be used by subsequent litigation. For example, individuals, trades unions and other commercial organisations may all seek redress if they feel that an incident has affected them in some material way. The statutes that govern the use of evidence vary from country to country. It is important that the personnel who are involved in incident investigations are familiar with at least the basic implications of these laws. For example, the following excerpt from the Law Commission for England and Wales provides an overview of criminal law in relation to the physical evidence and automated logs that can be gathered in the aftermath of an incident:

> "At present section 69 of the Police and Criminal Evidence Act 1984 requires a party to prove that the computer was working properly and was not being used improperly before computer evidence can be given. The Law Commission says this requirement is unnecessary and recommends its repeal... [The Law Commission's proposal on this point was implemented by section 60(1) Youth Justice and Criminal Evidence Act 1999 which provides that section 69 of the Police and Criminal Evidence Act shall cease to have effect.]
>
> The Commission recommends making automatically admissible those business documents which do not appear to be unreliable. At present all business documents are only admitted in evidence subject to the court's discretion. This discretion is exercised in different ways by different judges and magistrates, and parties cannot always predict whether the document will be admitted." [477]

The primary recipient of an incident report must first safeguard their system. They must then organise the acquisition of any evidence that might be relevant for the subsequent reconstruction and analysis of an adverse occurrence. This section has focussed on the acquisition of automatic logs and of physical evidence. The following section extends this analysis by looking in detail at the problems that arise when primary recipients must interview personnel in the aftermath of a safety-related incident .

6.2.2 Eye-Witness Statements

Witness statements are crucial to our understanding of the events that contribute to adverse occurrences. Without the evidence of those who were involved in an incident, it can be difficult or impossible to chart the ways in which multiple concurrent failures contribute to the eventual outcome. This data is particularly important for incidents that involve human factors issues. For example, the following excerpt from an incident report relies almost entirely upon the recollections of those involved. It is also instructive in that the analyst clearly does not take this evidence at face value:

> "At 0800, there were three persons on the bridge of Eternal Wind, the Mate, the 3rd Mate and the 4-8 AB[1]. The AB had been occupied writing up the deck log and plotting the position on the navigation chart, he had not been engaged in keeping a lookout after 0730 and, when interviewed, could not recall seeing any other vessels at all at that time. The Mate, who had been keeping his own lookout, at hand-over of the watch pointed out two vessels to the 3rd Mate, one northbound 13.5 miles to the west, the other four points on the starboard bow and southbound. Although Melina T would have been on the visible horizon of 8.5 miles at 0744 and had closed to a distance of four miles at 0800, the Mate had not seen the fishing vessel. The 3rd Mate, in taking over the watch, checked the horizon, using binoculars, and the radar, both on the 24 and 12 mile ranges, for other shipping. Visually he saw only the two ships handed over by the Mate, which he stated at interview were the only two targets indicated on the radar. He too did not see Melina T, which was at the same distance off as the southbound vessel and approximately midway between it and the ships head. Neither did he see the fishing vessel during the following 10 minutes, in which time it closed to a distance of 1.35 miles. It is evident that the lookout being kept aboard Eternal Wind was not effective. The 3rd Mate claimed that the reason for his not seeing Melina T was that he was blinded by the reflected glare of the sun. The strong glare was evident in a video film of the rescue, shot by one of the Eternal Wind crew-members, but despite this, the 3rd Mate was not wearing sunglasses." [520]

It can be difficult for the primary recipient of an incident report to determine the best time to interview the personnel who were involved in an adverse occurrence. If they meet with them in the immediate aftermath of an incident then feelings of shock and guilt can bias their responses. If they wait too long then memories of the incident may fade. There is also the danger that colleagues will gradually accept a shared view of events that may not initially have been held by all of the members in a group. Some organisations have established interview procedures to address these issues. For example, an initial debriefing session is held by the initial recipient. Subsequent interviews help to confirm the results of this preliminary meeting. They also help to elaborate any areas of remaining uncertainty. These subsequent interviews may be conducted by regional or national investigators or by the primary recipient depending on the seriousness of the occurrence [423].

There are many potential problems in conducting interviews. As we shall see, it is possible for the interviewer to bias responses by asking leading questions. For instance, asking 'why do you think the controller failed to spot this' presupposes that the controller actually did fail in the manner described. It is also possible to mis-interpret the responses that are provided to an interviewer. As a result, the US Occupational Health and Safety Administration (OSHA) has published a number of practical recommendations that are intended to guide interviews during incident investigation:

[1] This refers to an Able Seamen (AB) on the 4-8 watch

> “In general, experienced personnel should conduct interviews. If possible, the team assigned to this task should include an individual with a legal background. In conducting interviews, the team should: Appoint a speaker for the group. Get preliminary statements as soon as possible from all witnesses. Locate the position of each witness on a master chart (including the direction of view). Arrange for a convenient time and place to talk to each witness. Explain the purpose of the investigation (accident prevention) and put each witness at ease. Listen, let each witness speak freely, and be courteous and considerate. Take notes without distracting the witness. Use a tape recorder only with consent of the witness. Use sketches and diagrams to help the witness. Emphasize areas of direct observation. Label hearsay accordingly. Be sincere and do not argue with the witness. Record the exact words used by the witness to describe each observation. Do not ‘put words into a witness’ mouth’. Word each question carefully and be sure the witness understands. Identify the qualifications of each witness (name, address, occupation, years of experience, etc.) Supply each witness with a copy of his or her statements. Signed statements are desirable.” [649]

Such pragmatic advice may seem like common sense. It is surprising, however, that many incident reporting systems rely upon ad hoc interview techniques. It is important to provide more coherent support when different interviewers are used to gather information about incidents that are reported to regional, national and international systems. There is a danger that inconsistencies in the elicitation of interview data can introduce systematic biases in the causal analysis of adverse occurrences.

Interview Structures

When providing advice or drafting procedures to support interviews about adverse occurrences, there are a number of issues to consider. These are illustrated by the US Department of Justice’s guidelines of eliciting eye-witness statements:

> “When interviewing a witness, the preliminary investigating officer should:
>
> 1. Establish rapport with the witness.
> 2. Inquire about the witness condition.
> 3. Use open-ended questions (e.g., What can you tell me about the car?), augment with closed-ended questions (e.g., What colour was the car?). Avoid leading questions (e.g., Was the car red?).
> 4. Clarify the information received with the witness.
> 5. Document information obtained from the witness, including the witness identity, in a written report.
> 6. Encourage the witness to contact investigators with any further information.
> 7. Encourage the witness to avoid contact with the media or exposure to media accounts concerning the incident.
> 8. Instruct the witness to avoid discussing details of the incident with other potential witnesses.” [582]

These guidelines are intended to support interviews during criminal investigations. There are, however, a number of constraints that complicate their application to incident reporting. For instance, economic considerations may prevent face-to-face meetings if colleagues are geographically dispersed or if their work involves significant amounts of travel, as in the case of pilots. Face to face interviews can also compromise the confidentiality of a system if the other members of a team become aware of such meetings. There are a range of further issues. For instance, it is important to determine whether or not a predefined set of questions will be used to structure the course of an interview. Similarly, it is important to decide whether or not to focus respondents answers by providing a predefined set of responses:

1. *Unstructured or flexible interviews.* These typically have a set of predefined topics but no prescribed questions. These topics might include the interviewees observations about the state of the system in the run-up to the incident. The interviewee might be prompted to provide their opinion about causal and mitigating factors. They could also be asked about the ways in which an incident was detected. These topics help to identify generic areas of concern that are common to many different incident investigations. For instance, the New Zealand Department of Labour urges health and safety representatives to ask a number of questions. "*Who?* Get the names of everyone involved, near, present or aware of possible contributing factors. *What?* Describe materials and equipment involved, check for defects, get an exact description of chemicals involved, etc. *Where?* Describe exact location, note all relevant facts, i.e. Lighting, weather, etc. *When?* Note exact time, date and other factors, i.e. shift change, work cycle, break period, etc. *How?* Describe usual sequence of events and actual sequence of events before, during and after the accident. *Why?* Find all possible direct and indirect causes AND How to keep it from happening again." [654] The general nature of these questions leaves the interviewer free to phrase them in a form that is appropriate to the particular incident under investigation. The interviewer is free to follow the interviewees' replies and to find out personal opinions in response to previous answers. There are a number of dangers with this approach. In particular, interviewers can be 'seduced' into pursuing the ideas and recollections of articulate interviewees. There is also a danger that the interviewee can lead the interviewer into prolonged discussions about topics that have little significance for the overall understanding of the incident under investigation. Unfortunately, it can be extremely difficult to determine whether this is a deliberate intention or an innocent preoccupation of the interviewee [686].

2. *Structured interviews.* These rely upon a tightly defined set of questions that are, typically, asked in a predefined order. There is little scope for exploring individual attitudes. This approach is often used in the immediate aftermath of an incident when a primary recipient simply needs to gain a coherent overview of the occurrence. A more prolonged investigation of individual attitudes can either be postponed until more is known about an incident or can be incorporated into stress counselling. Structured interviews can also be used to ensure that the minimum set of information is gathered about relatively minor incidents. This is important if organisations are to meet the documentation requirements that are often specified by regulators for adverse occurrences. There are a number of limitations with this approach. In particular, it can be difficult to ensure that the minimum set of questions actually capture all of the relevant information about an incident. There is also evidence that individual interviewers can also bias answers to pre-defined questions. Such concerns potentially jeopardise some of the supposed benefits of this structured approaches over unstructured interviews [360].

3. *Semi-structured interviews.* In semi-structured interviews, the interviewer may have a list of pre-defined questions that they can draw upon during the course of an interview. Some of these questions might be omitted if they are considered not to be relevant to a particular incident. Other questions can be introduced if particular issues are raised during the interview. OSHA recognises that this approach is often inevitable given the diversity of incidents that can occur:

    > "Prior to the interviews, the team leaders and members shall develop key, critical and screening questions to ask all witnesses. Such questions may be written down and provided to all interviewers. While a specific list of questions is highly desirable, it may be more practical in some cases to have only a list of the topics to be covered. This list shall be developed before any interviews are conducted and shall include: 1 What is your name, address, telephone number, job, and employer? 2 How long have you done your present job? Have you ever seen any problem like this before? 3 Where were you at the time of the accident? What were you doing? Is that your normal job? Did you notice anything unusual? 4 How did you discover the accident? Were you close enough to physically sense (see, hear, feel, smell) anything?" [647].

In order to maintain consistency, several incident reporting systems distinguish between 'mandatory' questions that are designed to satisfy regulatory requirements for the documentation of an incident. Other questions are explicitly labelled as optional.

4. *Prompted interviews.* These are a particular form of semi-structured interview. They consist of a list of questions that are deliberately designed to provoke more detailed responses from the interviewee. For example, the interviewer may begin by asking; what exactly did you see? After an initial response they can then elicit further information by asking; can you tell me a little more about that? Alternatively, the user can be prompted to provide further explanation by asking; what do you mean by...? As with flexible interview techniques, there is a danger that the interviewee can deliberately lead the interviewer away from significant areas of investigation. There is also a danger that they will focus on hear-say rather than direct observations of an incident.

5. *Closed response interviews.* The previous types of interview technique have looked at the ways in which the interviewer asks questions of an interviewee. Other forms of interview focus on the ways in which an interviewee can answer those questions. For example, interviewers can ask interviewees to select their answer to a question from a number of cards that are laid out in front of them. Alternatively, preferences can be expressed by sorting the cards into a particular order. A more constrained version of this technique, relies upon asking the interviewee questions that can only elicit either yes or no as an answer. These approaches have the advantage that they place the interviewer in control of the course of the interview. However, they clearly restrict the interviewee's opportunity to express their opinions. Although these techniques have been exploited by market research organisations and in requirements engineering, they have not been widely used to support incident reporting.

Wellbank [858] observes that the more structured an interview, the greater the interviewer's control. As a result, greater skill and expertise is required if flexible or semi-structured techniques are to be used. Preece et al [686] comment that structured interviews also provide considerable benefits if interviewers must elicit information from domain specialists. There is a danger with more open-ended questions that the interviewer may not be able to interpret the technical information that this being provided in response to a particular question. This analysis has important implications for particular domains. For example, in air traffic control there is often the requirement that any interview procedures be conducted by controllers with at least ten years of experience in a particular centre [423]. However, in medicine it is certain that no individual will possess the complete range of technical skills that are necessary to understand the many different factors that contribute to particular incident. Even in the case of air traffic control, skilled controllers are unlikely to have the technical expertise to understand the complex hardware and software interactions that can contribute to systems failures.

It is important not to underestimate the costs of interviewing contributors and witnesses in national and international systems. For instance, the UK CIRAS rail reporting system sends a investigator out to conduct a follow-up interview in response to every report form that is submitted. Similarly, NASA personnel go back to the contributors of many ASRS submissions. This approach requires considerable resources. There must be enough trained analysts to elicit the necessary information during follow-up visits. Alternatively, novel computational techniques might be recruited to improve the quality of information that is initially contributed in response to an incident. These techniques might, therefore, reduce the expense associated with site visits. Equally importantly, they might also avoid the biases that affect follow-up interviews. A number of social concerns must affect contributors during safety-related discussions with external interviewers. Eliciting more information in the immediate aftermath of an incident also helps to reduce any delay between the contribution of a report and a follow-up interview.

The problems of extracting information from domain experts has been addressed by work on knowledge elicitation in general and by computer-aided interviewing techniques in particular [725]. These interviewing techniques, typically, rely upon frames or scripts that are selected in response to information from the user. For example, the user of an air traffic management system might first be prompted to provide information about the stage of flight in which an incident occurred. If it happened during landing then a script associated with that stage of flight would be selected. This might provide further prompts about the activities of arrivals and departures officers or about specific items of equipment, such as minimum safe altitude warning (MSAW) protection. These detailed questions would not be appropriate for incidents during other stages of flight, such as those filed during en route operations.

The relatively simple script-based techniques, described above, offer a number of further benefits. In particular, the use of computer assisted interviewing can reduce the biases that stem from the different approaches that are used by many interviewers. Inter-analyst reliability is a continuing concern in many incident report systems [414]. The scripts embodied in computer assisted interviewing systems might also be tailored to elicit particular information about regulatory concerns. For instance, if previous accidents had indicated growing problems with workload distribution during certain team-based activities then scripts could be devised to specifically elicit information about these potential problems. Of course, this analysis must be balanced against the obvious limitations of computer-based interviewing techniques [725] . Further evidence is needed to determine whether the weaknesses of computers assisted interviewing in employment selection or the analysis of consumer behavior also apply to their application in incident reporting.

Interview Formats

The structure of the questions and responses that are expected from an interview represent one of several issues that must be addressed by primary recipients. They must also decide upon the format of any elicitation exercises. There are a number of alternative approaches ranging from one-to-one interviews through to team meetings and focus groups. As before, the following comments also apply to investigators who follow-up these initial enquiries:

1. *Individual interviews (one to one).* This has the potential benefit of being relatively informal. Questions can be asked to clarify any of the information that was uncertain from the forms mentioned in Chapter 5. They can also be used to elicit information that might be missing in the original submission. This approach also has the benefit of protecting confidentiality and, as a result, has been recommended by several regulatory agencies: “Witness interviews shall always be conducted in private unless the witness requests otherwise” [647]. The problems are that the interview can be seen as combative and antagonistic if the interviewee lacks the support of their colleagues and workplace representatives. It is usually better to conduct interviews with two investigators present in the room and to allow the personnel involved to bring in a colleague or other representative.

2. *Interview panels (many to one).* This approach can avoid the inter-personal problems of a one-to-one interview. Several people, including friends and colleagues of the person being interviewed, can meet to discuss the occurrence. However, if such a meeting is not chaired correctly then it can appear to be an inquisition rather than a meeting to elicit necessary safety information.

3. *Team-based interviews (one to many).* In this approach, one interviewer meets with members of the shift during which an incident occurred. This reduces the inter-personal problems that can arise from a one-on-one interview. It may also help to uncover information from others who were present but not directly involved in an incident. The disadvantages include the practical problems of gathering everyone together but also the problems of accounting for group dynamics. The interview may be dominated by forceful personalities within the group. They may also compensate for the failures of one of their friends or exacerbate the weaknesses of those who are less popular.

4. *Group discussions (many to many).* This approach enables teams of investigators and works to get to together to discuss an occurrence. This has the benefit that neither group need be seen to be ‘in control’. Conversely, of course, it can lead to a general meeting that produces few tangible results and which reduces to a very general discussion.

There a number of techniques that primary recipients can exploit to address some of the problems that stem from team-based interviews. In particular, it is possible to use a number of map-based plans to illustrate the flow of conversation during a meeting. Figure 6.3 illustrates this approach. Firstly, an observer notes down the name and position of every person in the room. Secondly as each person contributes to the discussion, the observer draws a line between that person and their intended audience. At the end of each meeting these diagrams can be inspected to determine which of the participants contributed most to the meeting. If particular individuals are shown to have dominated proceedings then the interviewer must determine whether this reflects their involvement in the occurrence. If not then some of the findings from the meeting may have been biased by the views of this individual. If other people are shown not to have participated so actively in a discussion then follow-up interviews can be used to determine whether or not their views were adequately reflected during the course of the meeting. Such differences in participation can even out during the course of a meeting. It can often be helpful, therefore, to begin a new diagram each time the topic of conversation changes. This can reflect the way in which different individuals may have different degrees of participation in the lead-up to an incident and in any mitigating actions.

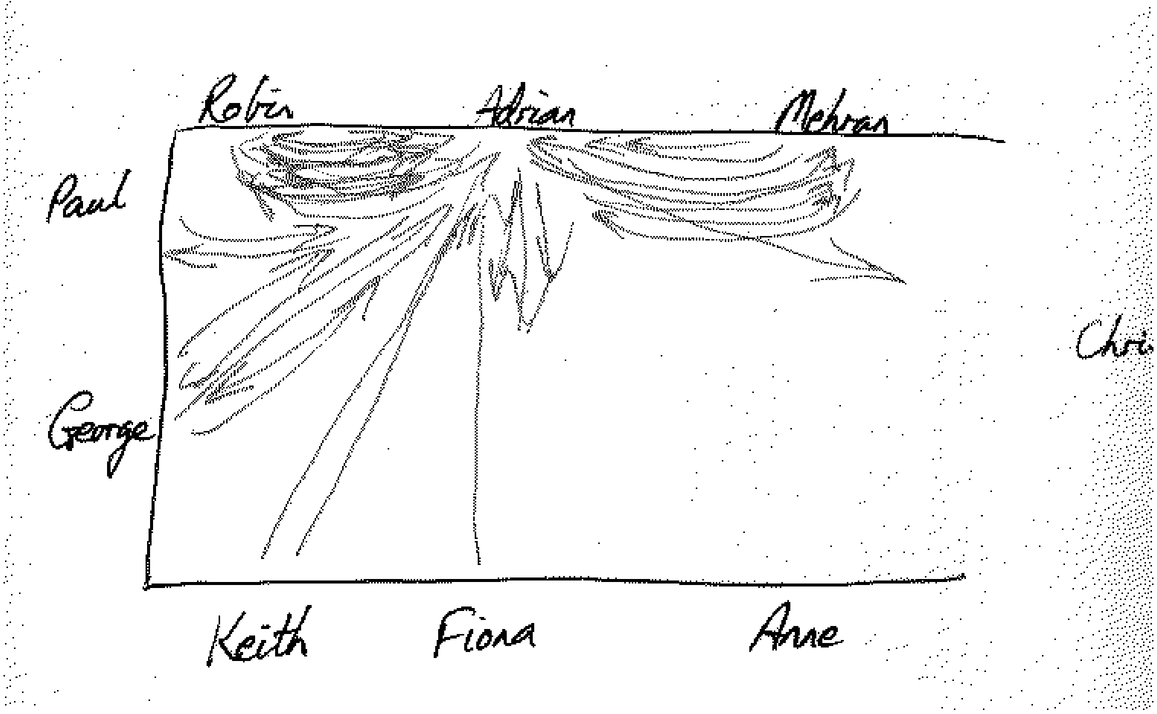

Figure 6.3: Interview Participation Diagram

This approach can also be used post hoc if the interviewees agree to have their contributions recorded. This raises a number of further issues. Audio tapes provide important reminders of passing comments that can easily be overlooked as interviewers struggle to control and direct a meeting. However, they lose the facial expressions, gestures and other forms of non-verbal communication that can be necessary in interpreting the force and meaning of an utterance [226]. Alternative, video recordings can provide much more of this contextual information. Unfortunately, our ability to analyse this data has not kept up with our ability to collect it. The rich information that can be obtained from such recordings makes it correspondingly more difficult to transcribe and analyse [724]. For both video and audio recordings, it is important to remember the OSHA directive that "interviews shall not be tape recorded as the only record of the interview" [647]. If such recording devices are used then the interviewer must also arrange for an alternative physical transcript in case the devices fail or the recordings are later corrupted.

There are a number of key principles that should guide any interview process. Firstly, the interview should have a purpose. As mentioned previously, interviews are costly in terms of the time needed to prepare for and attend such meetings. They also involve considerable resources if their results are to be accurately transcribed and analysed. Secondly, the results of any interview should be recorded in either written or electronic form so that both the interviewer and the interviewee can subsequently review the products of the meeting. Thirdly, these results should be reviewed. There is little point in conducting such an exercise if it is not to be used as part of a subsequent enquiry. Finally, the findings from any interview should be documented in a formal way and (ideally) communicated to the interviewee. Otherwise, such meetings can increase stress on an individual and ultimately lead to rumour and discontent within a working group.

Legal Issues Surrounding Eyewitness Statements

Previous sections have argued that even within no-blame systems, there are circumstances in which an initial investigation can uncover criminal actions. It is for this reason that OSHA recommend that each interview panel should include at least one member with at least some legal training [649]. The law governing witness statements varies from country to country, although there are a number of common features such as rules against hearsay. Hearsay, in a general sense, refers to the repetition of information received from others rather than from personal knowledge. Within the UK legal system there are a number of exceptions that make such information admissible in court. In particular, hearsay can be used for the purposes of identification. The Law Commission for England and Wales have recently sought to extend this exception:

> “...the identification exception extends only to identifications of people, and referred to cases such as Jones v Metcalfe (31) as revealing a deficiency in the law. Thus, where it is sought to establish the registration number of a car involved in an incident, and an eye-witness A, who saw the incident, related the number to B, who did not, it is inadmissible hearsay for B to tell the court what the number was for the purpose of proving which car was involved.” [477]

It is a sobering thought that many accident and incident reports make extensive use of hearsay evidence that would not be admissible in a court of law. The following extracts illustrate the complexity of legal provisions regarding eyewitness evidence. It describes a number of exceptions that apply to the rule of previous consistent statements. This is significant because under this rule when a witness does give evidence it is not usually possible to put in evidence previous statements by that witness. As a result, evidence gathered at interview is ‘superceded’ by the witness’ direct testimony. This raises particular problems for the subsequent handling of any incident enquiry if the the witness cannot significant information when it comes to trial. Previous statements cannot be used to reinforce the original terms of an identification or description.

> “(4) What we called in the consultation paper (5) the rule against previous consistent statements (and what others have called the rule against narrative) is the rule that such a statement cannot even be used to enhance the credibility of the witnesses oral evidence, by demonstrating the consistency of his or her story. This rule is subject to several exceptions.
>
> 10.88 A witness may be cross-examined on an oral or written statement made before the trial which is inconsistent with his or her oral testimony. The evidential use of the earlier statement is governed by the common law. If the witness accepts the earlier statement as being true, it is evidence of its facts; but where the witness denies the truth of the earlier statement it is not evidence, being nothing but hearsay, in which case the earlier statement reflects only on the witnesses credibility. If the witness does not admit making the earlier statement then the making of the statement may be proved.
>
> 10.63 A witness may refresh his or her memory from a statement in a document made contemporaneously with the events it concerns and while the facts were fresh in his or her memory. If the statement was recorded by someone else, the witness may nevertheless

> make use of it if the witness verified or adopted the statement. The document does not become an exhibit merely because a witness refreshes his or her memory from it."

The previous analysis focuses on criminal law within England and Wales. The intention is not to identify generic issues that affect all legal jurisdictions. In contrast, these provisions have been used to illustrate the importance of ensuring that primary recipients understand at least the basic legal framework that supports any subsequent litigation. If they do not have an appreciation of these constraints then any subsequent interpretation of the evidence may be open to legal challenge. These considerations affect confidential, proportionate-blame systems as much as they affect open reporting systems. For instance, interviewees often ask investigators about the legal implications of answering particular questions. It is important that the answers to such questions are both honest and truthful. It is also important to stress that no-blame systems continue to operate within the rules established by national legal systems.

Interpreting Eyewitness Statements

Previous sections have described several interview structures ranging from flexible question and answer sessions through to more restrictive closed response approaches. We have also introduced different interview formats including one-to-one reviews and many-to-many group meetings. Previous sections have also briefly described some of the legal issues, such as hearsay and the rule against previous consistent statements, that must be considered when gathering evidence about adverse occurrences. In contrast, this section looks more closely at the reliability of witness statements and the factors that can influence individual recollections of incidents and accidents.

There have been numerous experimental studies of eye-witness recollection [7, 223, 860]. A typical method involves showing a witness a simulated 'crime'. They are then asked if the 'criminal' is in a line-up potential suspects. If they are in the line-up then they are asked to identify them. Witnesses show a bias towards answering yes to the first of these questions irrespective of whether the criminal is actually in the line-up [223]. As Wickens notes; this would not be so worrying if individual eye-witness recall of brief incidents were not so poor [863]. He argues that studies into eye-witness responses reveal numerous biases that can affect both recognition and judgement. For example, individuals who express the greatest confidence in positive identifications are typically the least sensitive observers. Informing participants that a suspect may not be in a line-up can significantly reduce potential false-positives [223, 759]. They also argue that dressing individuals as similarly as possible will not only reduce the likelihood of biasing witnesses towards certain individuals but will also reduce the 'false alarm' rate. There are other factors that can bias individual eye-witness statements. For instance, Steblay identifies what has become known as the 'weapon focus' [758]. This biases the eye-witness to focus their attention on any weapon that is used in a crime rather than the perpetrator or the victim.

The basic psychological research into the eye-witness recollection of crimes has some relevance to accident and incident reporting. For example, it is possible to find evidence of the confidence bias in incident reports. Individuals who express the greatest confidence in their interpretation of an event may not be the most sensitive observers. This extension of the existing psychological literature is, partly, supported by judicial findings that must weigh the evidence provided by eye witness statements. For example, the following except shows how doubt can be cast on the evidence provided by witnesses who express undue confidence in their analysis. It is drawn from an OSHA case following an explosion in a detonator factory. The initial blast led to a secondary explosion involving a trailer that was parked nearby. The original judgement cleared the company of two violations of the US Occupational Safety and Health Act of 1970. The following quotation comes from a judicial review of the first decision and, therefore, reviews the quality of evidence provided by various witnesses:

> "I do not find the testimony of Prows and Del Regno summarily referred to by the majority to be compelling. First, neither Prows nor Del Regno testified that the cited trailer under the conditions existing at the time of citation was a service magazine. They offered only general opinion testimony to the effect that a trailer loaded with explosives

> and not moved for 'several days' or until 'ultimately loaded' would be a 'service magazine'. While neither Prows nor Del Regno gave further substantiation or qualification to the term 'several days,' I note that Prows made the contradictory statement that even 'one day would be too long'. Finally, Prows did not make a specific objection that the trailer was indeed in violation of the quantity-distance requirements during his prior inspection. Rather, Prows only observed that 'a loaded trailer would exceed' the limits. Given the lack of evidence regarding the amount of explosives on the trailer, the length of time the trailer remained at the dock is not relevant, even under the majority's test. In sum, there is no evidence that the trailer remained at the dock without fuses being loaded onto it and without proceeding to shipment. Therefore, I conclude that the trailer was spotted at the building for loading and shipping purposes rather than for the intermediate storage of explosives." [642]

This quotation is interesting because it provides indirect evidence to support the previous psychological studies into eye witness evidence. These studies identified a form of bias that occurs when over-confident witnesses are likely to miss significant information. The previous citation, arguably, shows that judges develop considerable expertise in spotting the flaws in evidence which is provided by such witnesses. However, such an interpretation goes well beyond the more focussed laboratory studies that characterise previous research in this area. More work is clearly need to determine whether or not these biases affect witness reports in the aftermath of incidents. Similarly, further research is needed to determine whether or not individual judges become skilled in filtering for these biases. For example, there is some evidence to suggest that the power of these effects varies even within the legal profession. Brigham and Wolfskiel surveyed 89 public defenders, 69 state prosecutors and 77 private defence attorneys in Florida [95]. 75% of prosecutors believed that witnesses who are more confident are more likely to be accurate. However, only 40% of defence attorneys agreed with this statement. It is readily apparent, however, that considerable weight is often placed upon the evidence of witnesses who recognise the limits of their statements. This is particularly apparent when reviewing the treatment of expert testimonies before the Occupational Safety and Health Review Commission that resolves disputes arising out of enforcement actions brought by the US Secretary of Labor:

> "We would comment that this was a difficult case, which we have decided solely on the preponderance of the evidence test. Weighing and reconciling conflicting opinion testimony from expert witnesses is never a simple task. Here, we were impressed by the candor of Professor Hochman, who did not attempt to convince us that the wires could not possibly have been broken before the accident. Instead, he explained that, because of the court's injunction, he was not able to perform the necessary examination in order to make that determination. He explained how, without such an examination, one kind of break may be mistaken for another. His testimony leads us to find that the other witnesses' opinions were formed without adequate empirical data to draw definitive conclusions." [645]

As mentioned, there has been relatively little work into the biases that affect eye-witness statements in the aftermath of incidents and accidents. Most previous research has focussed on individual and group recollections of criminal acts. These studies have been used to inform police procedures during the gathering of evidence for subsequent prosecutions. They have not been used primarily to inform safety improvements. As a result it is difficult to know whether or not observed behaviours can be used to help interpret witness statements in these two different domains. For example, it is possible to find parallels with the 'weapon focus' mentioned above. Eye-witness' who observe major equipment failures often focus on the behaviour of that equipment in subsequent accounts of an incident. As a result, they often omit important information about the behaviour of other systems or operators who indirectly influenced the eventual failure of that equipment. This analysis also has strong links to psychological research into 'post-event' reconstruction. This examines the ways in which individual memories change over time [500]. For example, an individual may be asked to observe a scene. They are then provided with information that is either consistent or inconsistent with the image that they have observed. Later when asked to recall aspects of that scene, the responses of individuals who received inconsistent information can be shown to be less reliable than those that had the reinforcement of consistent information. In psychological terms these studies are important because they long-term memory might be shaped by subsequent events. The legal implications of post-event reconstruction are clear [859]. For example, eye-witness evidence in the detonator explosions investigated by OSHA, cited previously, indicates the diversity of opinions that can exist over relatively straightforward estimates of physical distance even when supported by photographic evidence:

> "At the time of the explosion on the production line, a semi-trailer truck was parked at the loading dock adjoining the work bay. Referring to a photograph in evidence, Harrold Owen, Respondent's president, testified that the distance between the end of the loading dock and the work bay was 48 feet. Other witnesses estimated the distance as 10 feet and 20 feet." [642]

These effects need not, however, simply be seen as the effects of post-event reconstruction. They can be interpreted as the result of social influences rather than more direct cognitive effects. For example, the relative distances cited in the previous excerpt were used in a more complex argument about the safe positioning of the trailor. The witnesses were not, therefore, simply recalling a physical distance. They were providing evidence that, in turn, supported or weakened particular lines of legal argument. Hence their recollections might have been influenced by their knowledge of the context in which their evidence was being elicited.

There remains considerable disagreement about the impact of repressed memory syndrome on eye witness testimony. As with previous studies, most of the work in this area has not focussed on eye-witness statements in the aftermath of incidents and accidents [861]. It has, in recent years, focussed on recollections of childhood abuse. Critics of this work have shown that "children who witness traumatic events seem to have trouble forgetting it rather than showing signs of repression" [859]. However, Lindsay and Read have also shown that false autobiographical memories can be created by suggestion and by repeated imagination [495]. They can also be correlated with a belief in the concepts of repression and recovery of repressed memories and by hypnosis or hypnotic-like interventions.

Cultural Issues

The previous paragraphs have briefly reviewed the many complex factors that must be considered when interpreting eyewitness statements. For example, we have cited studies in which individual recollections of an incident can be affected by prompts and questions that they receive during post-event reconstruction. These factors have received considerable attention as a result of the increasing number of unsafe convictions in which DNA tests have been used to exonerate individuals who have been convicted on the strength of eyewitness statements [180]. As a result, national guidelines have been developed to minimise such influences during subsequent interviews [861]. For instance, the following excerpt provides the US Department of Justice's guidance on the interpretation of eye witness testimony:

> "*Principle:* Point-by-point consideration of a statement may enable judgement on which components of the statement are most accurate. This is necessary because each piece of information recalled by the witness may be remembered independently of other elements. *Policy:* The investigator shall review the individual elements of the witness statement to determine the accuracy of each point. *Procedure:* After conducting the interview, the investigator should:
>
> 1. Consider each individual component of the witness statement separately.
> 2. Review each element of the witness statement in the context of the entire statement. Look for inconsistencies within the statement.
> 3. Review each element of the statement in the context of evidence known to the investigator from other sources (e.g., other witnesses statements, physical evidence)."[582]

There are further issues that arise in using witness statements from individuals who have been trained within particular organisational cultures. Again, many of these concerns stem from the use of evidence in police investigations. However, the underlying issues also affect the use of witness statements in more general investigations. For instance, police officers have often been criticised as witnesses in criminal cases because they may hold certain beliefs and biases that affect their perception, recognition and recall of events in a way that might not affect other members of the public. These biases can stem from the recruitment and selection process, from training, from working culture or from experience. For example, training manuals have in the past directed officers to look for particular characteristics of groups. The clothes that they wear, the way in which they stand and walk, their use of language all provide indications of a potential criminal intent. This reinforces stereotypical categories that can support everyday police tasks. These categories can also reinforce inappropriate cultural stereotypes that lead individual officers to ill-considered assumptions about the perpetrators and course of a crime. In the UK, these concerns crystalised in the Macpherson's Inquiry into the death of Stephen Lawrence [511]. Stephen Lawrence was murdered by a group of five or six white youths while he waited for a bus on 22nd April 1993 . Initially it was thought that he had been involved in a fight rather than an unprovoked racist attack. The subsequent investigation failed to result in the conviction of anyone involved in the incident. Prolonged police investigations, in two distinct phases, produced only one witness. The Police Complaints Authority engaged the Kent Police to investigate complaints by Stephen Lawrence's parents that the first Metropolitan Police Service (MPS) investigation had been bungled. The resulting report roundly criticised many aspects of the MPS investigation. Public concern over the findings of this document and the justified indignation of Stephen Lawrence's parents led the Home Secretary to instigate a more general inquiry. The resulting Macpherson report proposed the following definition for 'institutional racism':

> "*Institutional Racism* consists of the collective failure of an organisation to provide an appropriate and professional service to people because of their colour, culture or ethnic origin. It can be seen or detected in processes, attitudes and behaviour which amount to discrimination through unwitting prejudice, ignorance, thoughtlessness, and racist stereotyping which disadvantage minority ethnic people." [511]

It is readily apparent that the Macpherson report deals with the failure of a criminal investigation rather than a 'near-miss' incident. However, the findings of this inquiry are extremely important for any reporting system that collects and analyses accounts of complex human behaviour. Organisational factors not only effect the sorts of occurrences that are contributed, through its reporting culture, they also affect the organisations interpretation and response to those occurrences. The problem of institutional racism, or other forms of discrimination, are clearly not restricted to the UK police service. The Macpherson report goes on to describe in precise detail how the problem of institutional racism affected many different stages of the investigation into Stephen's death. For instance, the initial investigations failed to consider the evidence of the main witness that Stephen Lawrence had been the victim of an unprovoked attack. This inquiry is unusual in that it provides arguably the only analysis of the corrosive effect that organisational 'bias' has upon a professional organisation. The concern is that if these factors affected the Metropolitan Police's investigation of a murder then the biases may be even more pronounced in the elicitation and analysis of evidence in less serious incidents by less well-trained personnel [379]:

1. "Inspector Groves' insensitive and racist stereotypical behaviour at the scene. He assumed that there had been a fight. He wholly failed to assess Duwayne Brooks as a primary victim. He failed thus to take advantage of the help which Mr Brooks could have given. His conduct in going to the Welcome Inn and failing to direct proper searches was conditioned by his wrong and insensitive appreciation and conclusions.

2. Family Liaison. Inspector Little's conduct at the hospital, and the whole history of later liaison was marred by the patronising and thoughtless approach of the officers involved. The treatment of Mr and Mrs Lawrence was collective, in the sense that officers from the team and those controlling or supervising them together failed to ensure that Mr and Mrs Lawrence were dealt

with and looked after according to their needs. The officers detailed to be family liaison officers, Detective Sergeant Bevan and Detective Constable Holden, had (as Mrs Lawrence accepted) good intentions, yet they offended Mr and Mrs Lawrence by questioning those present in their house as to their identity, and by failing to realise how their approach to Mr and Mrs Lawrence might be both upsetting and thoughtless.

3. This sad failure was never appreciated and corrected by senior officers, in particular Mr Weeden, who in his turn tended to blame Mr and Mrs Lawrence and their solicitor for the failure of family liaison. The failure was compounded by Mr Barker in his Review.

4. Mr Brooks was by some officers side-lined and ignored, because of racist stereotyping particularly at the scene and the hospital. He was never properly treated as a victim (Chapter 5).

5. At least five officers, DS Davidson, DC Budgen, DC Chase, DS Bevan and DC Holden simply refused to accept that this was purely a racist murder. This (as we point out in the text) must have skewed their approach to their work (Chapter 19).

6. DS Flook allowed untrue statements about Mr and Mrs Lawrence and Mr Khan to appear in his statement to Kent. Such hostility resulted from unquestioning acceptance and repetition of negative views as to demands for information which Mr and Mrs Lawrence were fully entitled to make. DS Flook's attitude influenced the work which he did (Chapter 16).

7. The use of inappropriate and offensive language. Racism awareness training was almost non-existent at every level." [511]

Previous paragraphs have used the Stephen Lawrence inquiry to illustrate the ways in which cultural norms can bias the direction of police investigations. Whist the problems of institutional racism have not been identified to the same degree in other safety-critical professions, including medicine and aviation, it is possible to find other forms of organisational bias [409]. For example, many professional groups can influence the reporting behaviour of its members by exerting a strong normalising influence [342]. The esoteric nature of the knowledge and skills that are required by professions, typically, implies that their members are self-regulating. This affords a degree of protection from the general public. In exchange the members of the profession accept the 'social control' of their peers. This normalising influence is not common to all professions. For example, the role of the external regulator in aviation makes it more difficult to preserve the internal regulation of an 'old boy network'. However, there are other forms of profession bias. In particular, the 'self-concept' has been used to describe the self evaluations that people make with reference to other groups of their peers. There are striking parallels between this analysis of the cultural barriers to professional change within the medical and aviation communities and the problems faced by the Metropolitan Police in the aftermath of the Macpherson report:

> "Since work is the central aspect of being for many, the internalised values of professional culture are likely to be important components of the self-concept. The positive aspects of professional culture, including prestige, contribute to a positive self-concept in the work domain and to self-esteem. Unfortunately, the negative aspects of the culture including the sense of invulnerability, also become integral parts of the self-concept. One of the more provocative findings regarding the self-concept is that individuals seek to maintain their established self-concepts, even when they are recognised as negative. The resistance of self-concepts to discomfirming evidence can explain why attitudes about personal limitations seem to fall on death ears and why change proceeds at a slow pace" [342]

This section has introduced a number of factors that complicate the elicitation and the interpretation of evidence from eye-witnesses. Some of these problems stem from basic properties of human cognition. For instance, it seems likely that individual memories of complex events can be affected by the witness' subsequent re-appraisals of the events they have observed. Other problems relate more narrowly to the biases that affect those individuals who collect eye-witness statements. It is relatively easy to guide evidence by posing leading questions or by suggesting particular lines of argument. Later sections have gone beyond the effects of individual bias to look at the cultural norms that prevent, or conversely promote, the effective use of eye-witness statements.

6.3 Drafting A Preliminary Report

A number of national and international bodies require that incident information is disseminated to other organisations that might be involved in similar adverse occurrences. For instance, the ICAO specify that if incident reports help a State to identify safety matters that are considered to be "of interest" to other States then that State should forward the information to them "as soon as possible". They are require that member States "promote the establishment of safety information sharing networks" that facilitate the free exchange of information on actual and potential safety deficiencies [384]. As a result, they require that member states should draft a preliminary report within thirty days of a severe incident and "as soon as reasonably practicable' for minor occurrences. For more severe incidents, the report must be sent to the State of registry of an aircraft or the State in which the incident occurred. It should also be sent to the State of the operator, the State of design and the State of manufacture. A copy of this prelimiary report must also be sent to states that provided relevant information, significant facilities, or experts. A copy must also be sent to the ICAO. For less sever incidents, the distribution requirements for a preliminary report are more limited:

> "The State conducting the investigation should upon request provide other States with pertinent information additional to that made in the Accident/Incident Data report.

Aviation is not the only domain in preliminary initial reports are used to warn other organisations about adverse occurrences. For example, the FDA require what is known as a 5-day report after the notification of a medical incident to a device manufacturer [258]. The International Atomic Energy Authority (IAEA) require a "short preliminary report" within one month of a nuclear incident being reported in a national incident reporting system coordinator [382]. Although there are significant differences in the regulatory requirements for these initial reports, there are also a number of common features. For example, the primary recipient of an incident report is often left to draft the preliminary report into less severe incidents. They must collate the available evidence in the manner described in previous sections. The primary recipient then use this evidence to perform an initial severity assessment. They typically, conduct an informal causal analysis of the events that contributed to the failure. The preliminary report is then passed to regional or national safety managers who can supplement the report if necessary. For more severe incidents, the task of drafting a preliminary report is typically to professional incident investigators.

6.3.1 Organisational and Managerial Barriers

Irrespective of who produced the initial report, safety managers must decide who should receive copies of this document. A number of factors influence their decision. Most importantly, managers must determine whether there is a significant risk of a similar incident recurring at other sites both inside and outside their organisation. If the preliminary report suggests that such a risk exists then information must be passed on. There are clear ethical and legal implications about any failure to pass on reports of previous failures if a similar incident does occur in the future. The decision to pass on a preliminary report can also be influenced by explicit requests to receive information on particular topics. For example, the European Turbulent Wake incident reporting system registered an interest in hearing about any of these incidents that involved commercial aircraft [547]. At a local level, managers may decide to pass on preliminary reports if they identify an incident as part of a regional trend. This depends upon a careful monitoring of incidents over time and, in the eraly stages of an investigation it may be impossible to accurately determine whether a particular occurrence does or does not form part of a wider pattern.

It may at first sight appear that preliminary reports should, by default, be broadcast as widely as possible. For instance, the International Atomic Energy Authority reporting system encourages national coordinators to provide information about all incidents that might be of international interest. In all cases, preliminary reports are followed-up by the publication of a final report:

> "Each participating member country designates a national Incident Reporting System (IRS) co-ordinator. An event report is submitted to IRS when the event is considered by the national co-ordinator to be of international interest. IRS when the event is considered by the national co-ordinator to be of international interest. Only events of safety significance are reported. When information is considered time sensitive, a short preliminary report is distributed within one month of the event." [382]

However, the decision to publish all preliminary information is not as simple as it might seem. Confidence in reporting systems can be jeopardised if a large number of preliminary reports are subsequently revised in the light of more detailed investigations. Warnings about potential incidents can threaten long-term safety if organisations forget to revise their initial corrective actions in the light of any subsequent findings. There is a danger that a large number of 'spurious' reports can mask preliminary information about more critical incidents. As a result, some incident reporting systems actively prioritise or filter the dissemination of these initial reports [806]. Only the most critical documents are released until more evidence is obtained about the causes and consequences of an adverse occurrence. Other systems adopt a multi-tier approach in which a succession of regional, national and international committees determine whether information about an incident should be passed onto the next level of investigation. This approach characterises some aspects of the European Space Agencies Alert system:

> "The providers and users of the information channelled through the European Space Agency (ESA) Alert System are the participating organisations. They play a key role in actively notifying failures and problems, which they do by initiating a PAI (preliminary alert information); they also participate in the investigation of a PAI. If the PAI is officially adopted it achieves the status of an ESA Alert. Participating organisations also act upon the information promulgated through an ESA Alert and provide feedback on the effectiveness of the suggested corrective actions. Each participating organisation nominates an Alert Coordinator who manages communications with ESA. Due to the sensitive nature of the information contained in an ESA Alert, ESA requires that all PAIs be subject to a rigorous scrutiny and a well defined authority is maintained for the release of an ESA Alert. The parties involved in these processes are: the ESA Alert Committee; the ESA Alert Focal Point; technical specialists. The ESA Alert Committee, chaired by the Head of the Product Assurance & Safety Department, ESTEC (Research and Development Arm of ESA), has overall responsibility to decide whether or not an identified failure or problem should be published as an ESA Alert. The ESA Alert Focal Point, is a centralised function within the ESA Product Assurance and Safety Department which administrates the ESA Alert system and maintains its effective functioning." [230]

There are further managerial and organisational factors that complicate the dissemination of initial information about incidents and accidents. There is a natural reluctance to publicise a potential failure in safety mechanisms prior to more detailed investigations. For example, the UK Major Hazard Incidents Data Service (MHIDAS)deliberately delays the publication of some incident reports in order to ensure that the information which it provides is as complete and as accurate as possible:

> "The database is updated every quarter, but incidents are not generally entered onto the database until a year after they have occurred so that as much information as possible can be collected for each incident from a number of different types of journals. Because of their nature, information published in reports soon after an incident occurred may be incomplete and for major incidents some early reports may contradict each other as the exact number of fatalities or injuries may not immediately be apparent. It is thus

> important that information on an incident is collected from as many information sources as possible." [323]

This quotation illustrates two different approaches to the publication of incident information. On the one hand, there is a requirement to provide as much accurate information about adverse occurrences as possible. This helps to ensure that lessons from past failures are not propagated into the future design and operation of safety-critical systems. On the other hand, there is a more immediate requirement to warn other operators about the potential for previous failures to recur. Clearly, there must be some alternate means of ensuring that adverse occurrences cannot be repeated in the twelve months before they appear in the MHIDAS database. This implies not only that preliminary reports must be published but also that thy must contain subsequent information for other organisations to be able to act on the warning. This is illustrated by the FDA's criticisms of Atomic Energy of Canada Limited's (AECL) response to exposure incidents involving the THERAC-25 linear accelerator. The following paragraph forms part of the FDA's response to a letter that was sent by AECL to each Therac user recommending a temporary 'fix' to the machine that would allow them to continue to be used:

> "We have reviewed [AECL's] April 15 (1986) letter to purchasers and have concluded that it does not satisfy the requirements for notification to purchasers of a defect in an electronic product. Specifically, it does not describe the defect nor the hazards associated with it. The letter does not provide any reason for disabling the cursor key and the tone is not commensurate with the urgency for doing so. In fact, the letter implies the inconvenience to the operator outweighs the need to disable the key. we request that you immediately re-notify purchasers." (FDA to AECL, Director of Compliance, centre for Devices and Radiological Health, cited in [486]).

This quotation illustrates how regulators will intervene if they believe that preliminary reports do not provide sufficient information about the potential risks of future failures. Such responses are usually symptomatic of a deeper breakdown in the relationship between the manufacturer or supplier and the organisations who must intervene to ensure the safety of the market place.

6.3.2 Technological Support

AECL's letter was sent almost twelve months after the initial incidents took place but less than one month after a lawsuit was issued by the first patient. These is evidence that this delay in issuing a preliminary report stems from the lack of any mechanism within AECL to follow-up on suspected accident or incident reports [486].Many organisations, therefore, explicitly publish deadlines for their initial response to an adverse occurrence [806]. Previous paragraphs have mentioned the FDA's 5 day rule and the ICAO's 30 day deadline for preliminary reports into the most serious incidents. Other organisations are forced to specify deadlines that vary according to the operational demands upon its staff. Figure 6.4 provides an extreme example of this. It illustrates the US Army's time-scales for the submission of preliminary, interim and final reports [806]. Unit commanders and safety officers can provide prelimiary reports over the telephone. The AGAR form, typically, provides an abridged form of interim information about accidents and incidents. The DA 285 form provides greater detail and in many cases represents a final accident report. IAI refers to an Installation Accident/incident Investigation, CAI refers to a Centralised Accident Investigation. As can be seen, these time-scales depend not simply upon the severity of the incident but also upon whether or not the unit reporting the incident is involved in a combat operation. For example, the abridged AGAR form can be used used under combat for category A and B accidents that would normally require the more exhaustive DA 285.

These deadlines can impose considerable burdens upon operational staff. As a result, organisations such as the US Army make extensive use of telephone notification procedures. Again, as can be seen from Figure 6.4 these are reserved for the preliminary reports associated with high-criticality incidents and accidents. The preliminary reports associated with less 'severe' incidents are submitted using the AGAR forms. Figure 6.5 shows the forms that operators must complete when they

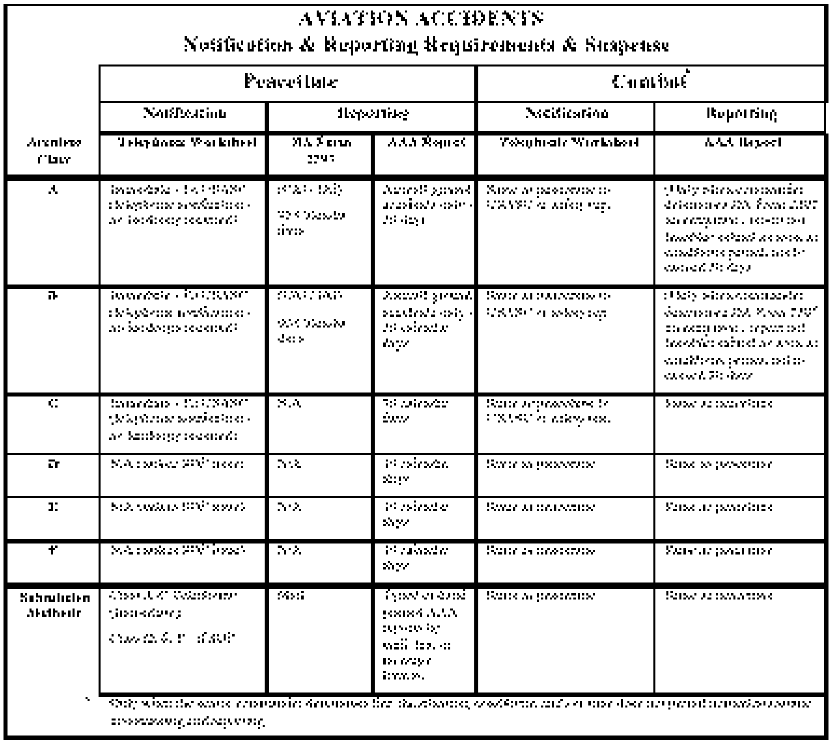

AVIATION ACCIDENTS
Notification & Reporting Requirements & Suspense

	Peacetime			Combat*	
	Notification	Reporting		Notification	Reporting
Accident Class	Telephonic Worksheet	DA Form 2397	AAA Report	Telephonic Worksheet	AAA Report
A	Immediate - To USASC telephonic notification - no backup required	[illegible] 90 calendar days	[illegible] 30 days	Same as peacetime to USASC or safety rep.	[illegible]
B	Immediate - To USASC telephonic notification - no backup required	[illegible] 90 calendar days	[illegible] 30 calendar days	Same as peacetime to USASC or safety rep.	[illegible]
C	Immediate - To USASC telephonic notification - no backup required	N/A	30 calendar days	Same as peacetime to USASC or safety rep.	Same as peacetime
D	[illegible]	N/A	30 calendar days	Same as peacetime	Same as peacetime
E	[illegible]	N/A	30 calendar days	Same as peacetime	Same as peacetime
F	[illegible]	N/A	30 calendar days	Same as peacetime	Same as peacetime
Submission Methods	[illegible]	Mail	[illegible]	Same as peacetime	Same as peacetime

* [illegible]

Figure 6.4: US Army Incident/Accident Reporting Procedures

receive a preliminary oral report of an incident. In passing it is worth noting that this form is quite different from some of those shown in Chapter 5. It is not intended to be completed by the staff who were involved in an incident nor is it expected that those staff would telephone-in an account of an incident. Instead this form represents a preliminary report because it is assumed that it will provide a record of the initial observations made by either a unit commander or by a trained safety officer. The degree of planning reflected in Figures 6.5 and 6.4 contrasts sharply with the FDA's criticisms of AECL. It also illustrates what is required if large, relatively complex organisations are to meet relatively tight deadlines for the investigation and analysis of adverse occurrences.

Information technology is increasingly being recruited to support more traditional communication media in order to meet the deadlines shown in Figure 6.4. For instance, Figure 6.6 illustrates the web-based interface to the US Army Aviation and Missile command preliminary incident reporting system [821]. As with the telephone form shown previously, this interface provides a rapid means for primary recipients to provide a safety management group with a preliminary report about an incident. The relatively open format, typified by the field labelled 'Description of incident', can be contrasted with the more tightly defined fields of the CIRS form illustrated in Chapter 5 [756]. This web-based system elicited reports directly from anaesthetists. The user of the CIRS system selects the types of incident from a predefined list of possible events. In contrast, the Army system covers may different engineering and military applications. As a result, the open field format provides greater scope for the initial analysis in the preliminary incident report.

6.3.3 Links to Subsequent Analysis

The previous sections have focussed on a number of organisational issues that complicate the dissemination of information contained in preliminary reports. It has been argued that concerns over the sensitive nature of information about system failure must usually be addressed by regulatory

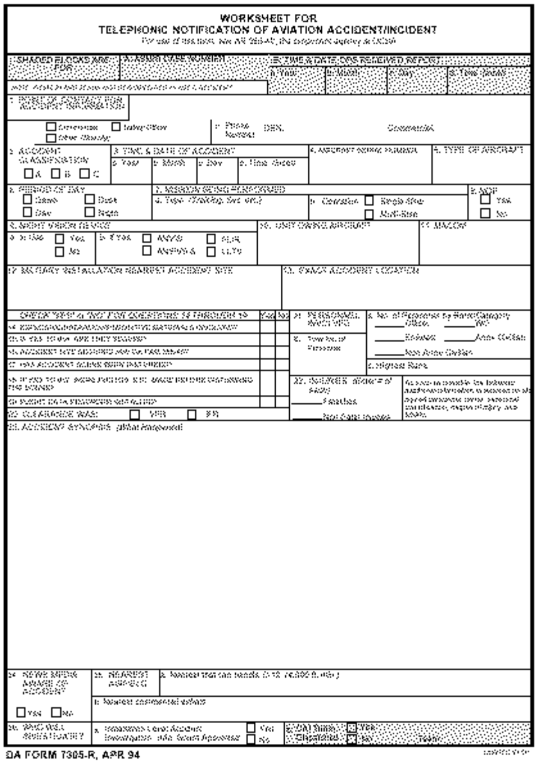

WORKSHEET FOR
TELEPHONIC NOTIFICATION OF AVIATION ACCIDENT/INCIDENT
[illegible]

SHADED BLOCKS ARE FOR [illegible]
A. [illegible] CASE NUMBER
B. TIME & DATE [illegible] RECEIVED REPORT
a. Year b. Month c. Day d. Time (local)

1. POINT OF CONTACT FOR ACCIDENT INFORMATION
2. ACCIDENT CLASSIFICATION □ A □ B □ C
3. TIME & DATE OF ACCIDENT a. Year b. Month c. Day d. Time (local)
4. AIRCRAFT SERIAL NUMBER
5. TYPE OF AIRCRAFT
6. PERIOD OF DAY □ Dawn □ Dusk □ Day □ Night
7. MISSION BEING PERFORMED
8. NOE □ Yes □ No
9. NIGHT VISION DEVICE
10. UNIT OWNING AIRCRAFT
11. MACOM
12. MILITARY INSTALLATION NEAREST ACCIDENT SITE
13. EXACT ACCIDENT LOCATION

CHECK "YES" or "NO" FOR QUESTIONS 14 THROUGH 19

21. PERSONNEL INJURED
b. No. of Personnel by Rank/Category
____ Officer ____ WO
____ Enlisted ____ Army Civilian
____ Non Army Civilian
c. Highest Rank

22. INJURIES
____ Fatalities

20. CLEARANCE WAS: □ VFR □ IFR

23. ACCIDENT SYNOPSIS (What Happened)

24. NEWS MEDIA AWARE OF ACCIDENT □ Yes □ No
25. NEAREST AIRFIELD
b. Nearest commercial airfield
26. WHO WAS NOTIFIED?

DA FORM 7305-R, APR 94

Figure 6.5: US Army Preliminary Incident/Accident Telephone Reports

intervention. Subsequent sections went on to briefly describe how fixed time-scales are usually imposed for the completion of preliminary and interim reports, such as the AGAR forms used by the US Army. Telephone procedures can be used to ensure that necessary information is passed from the primary receiver to central safety managers. Web-based systems are also playing an increasing role in the communication of initial information about adverse occurrences.

It is important to emphasise, however, that the drafting of a preliminary report only represents an initial step in the response to a safety-critical incident. This point is illustrated by the FDR's reporting system for the manufacturers of medical devices.

> "There are five types of Medical Device Reporting (MDR) reports that FDA requires the manufacturer to submit. Each type of report is to be submitted within the mandatory time frame by completing the appropriate form. MDR reports for manufacturers include a:
>
> 1. 30-day report,
> 2. 5-day report,
> 3. baseline report,
> 4. supplemental report, and
> 5. annual certification." [258]

The 30-day, 5-day and baseline reports represent refinements on the general concept of a preliminary report that has been presented in this chapter. If a manufacturer receives information about an MDR reportable event, they must submit a 5-day form within five work days after: (1) becoming aware that a reportable event necessitates remedial action to prevent an unreasonable risk of substantial harm to public health; or (2) becoming aware of an MDR reportable event from which FDA has made a written request for the submission of a 5-day report involving a particular type of medical device

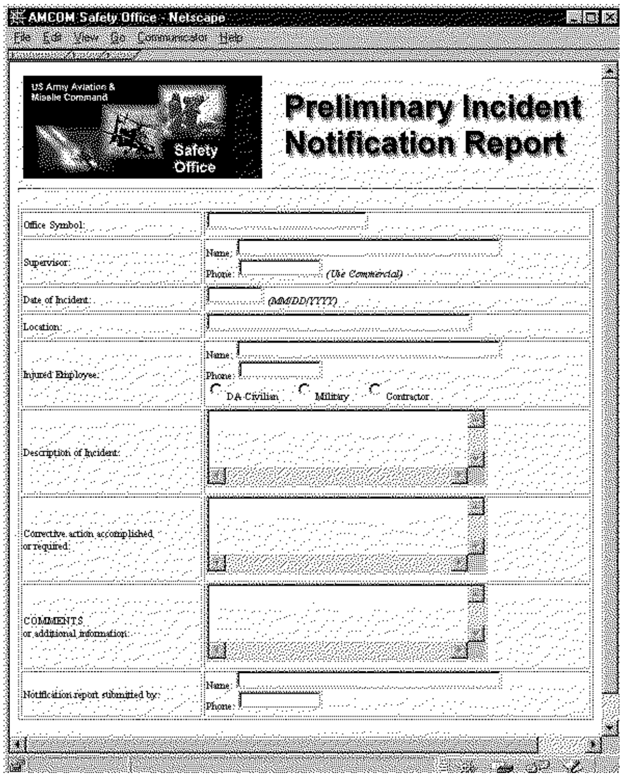
AMCOM Safety Office - Netscape
File Edit View Go Communicator Help
US Army Aviation & Missile Command
Safety Office

Preliminary Incident Notification Report

Office Symbol:	
Supervisor:	Name: Phone: (Use Commercial)
Date of Incident:	(MM/DD/YYYY)
Location:	
Injured Employee:	Name: Phone: ○ DA Civilian ○ Military ○ Contractor
Description of Incident:	
Corrective action accomplished or required:	
COMMENTS or additional information:	
Notification report submitted by:	Name: Phone:

Figure 6.6: US Army Aviation and Missile Command Preliminary Incident Form

or type of event. The thirty day report must be submitted by any manufacturer within 30 calendar days after becoming aware of a reportable death, serious injury, or malfunction. Baseline reports illustrate a further development of the preliminary report; they must be submitted in response to the first MDR reportable incident involving a particular device. This report provides basic device identification information including: brand name, device family designation, model number, catalogue number and any other device identification number. This information helps ensure clear, unambiguous device identification.

The last two classes of document required by the FDA's MDR scheme illustrate the way in which preliminary reports form part of a more complex process in which regulators may intervene to monitor any subsequent analysis, to oversee the implementation of any further remedical actions and to assess the overall effectiveness of those actions. Manufacturers must submit a supplemental report if they obtain additional information denoted as unknown or not available at the time of the preliminary 30 and 5-day reports. A supplemental report is also required when new facts prompt the manufacturer to alter any information submitted in the original MDR report. This must be submitted within one month of the receipt of the information.

Follow-up reports document important stages in the investigative process after the primary recipient has filed an initial notification with the MDR system. Typically, medical device manufacturers must seek this additional information by follow-up interviews with the end-users of their devices. This raises the question of how many attempts must manufacturers make to obtain additional contextual information about particular incidents. The FDA requires that a 'good faith effort' be made to obtain information. At least one request for information should be made in writing. In a sense, therefore, the preliminary 5 and 30-day reports help to identify the more detailed information needs that must be addressed during a subsequent investigation.

Annual reports provide a further monitoring tool for the FDA and the operators of the MDR system. Section 510(d) of the Federal Food, Drug, and Cosmetic Act (the act) [21 U.S.C. 360I(d)] provides that each manufacturer, importer, and distributor shall certify that they filed a certain number of medical device reports (MDR's) in the previous twelve months or that they did not file any MDR's. The legal requirement helps to ensure that the FDA keeps an overview of the relative performance of particular commercial organisations from year to year. By requiring that named individuals sign these annual reports, there is an additional means of verifying the internal MDR audit mechanisms.

The MDR procedures illustrate how prelimiary reports, at 5 and 30 days, can be used to provide an initial notification of an adverse occurrence. Previous sections have argued that these initial reports often contain omissions and inaccuracies. The FDA have addressed these concerns by providing for supplementary reports that are intended to resolve any ambiguities or gaps that could not satisfactorily be explained within the relevant time limits. Each of these reports, in turn, must be accounted for in an annual report that provides an overview of the longer-term safety record of an organisation. Importantly, this mechanism also forces individuals to document the multiple 5 and 30 day reports that can arise when the same device generates numerous incidents. Base line reports provide the necessary identification information to ensure that reports of these failures are not disguised by arbitrary distinctions within a product line. The key point behind all of this is that preliminary reports only provide an initial glimpse of the information that must be collected for more serious incidents. There must be some mechanism for ensuring that these additional details are collected and recorded. There must also be some means of assessing the effectiveness of the entire reporting process, for instance through annual surveys of incidents and accidents.

6.4 Summary

This chapter focussed on the responsibilities of the 'primary recipient'. This term is used to describe the supervisors, managers or other nominated personnel who first receive an incident report. Initially, their first priority is to safeguard their system. This can involve removing operators from positions of control if their involvement in an incident makes them susceptible to further 'errors'. It may also force them to instigate back-up procedures or to restrict the level of service that is provided. It is critical, however, that any remedial actions should not exacerbate the consequences of any initial failure. A number of factors were identified that can combine to frustrate attempts to safeguard the system. These include poor training in emergency procedures and a lack of situation awareness that can prevent primary recipients from accurately predicting the consequences of any intervention. Their tasks can also be complicated by time pressures in the aftermath of an incident. Lack of information and a lack of necessary system support can deprive primary recipients of the necessary resources to effectively direct their interventions. The pressing need to preserve levels of service, for example in air traffic control, can also further complicate attempts to safeguard a system. Previous sections then went on to review a number of emergency management procedures that can be used to address many of these potential pitfalls. Documented procedures, reinforced through simulated emergency training, have proven to be effective in many different domains. There are, of course, concerns that such techniques may do little more than establish stereotypical responses that can even hinder an individual's ability to respond to pathological failures. One solution to this potential weakness is to ensure a close link between the scenarios that are used during simulated emergencies and the incident information that is gather by reporting systems in similar organisations.

Later sections went on to discuss the problems that primary recipients face in gathering automated data about adverse occurrences. It can be difficult to predict which logs will actually contribute most to any subsequent investigation. In consequence, many regulatory organisations specify a minimum list of information sources that must be secured after any incident. It is important, however, to realise that many automated systems cannot be relied upon to produce accurate information about a failure. For example, the loop recording facilities of cockpit voice recorders make it particularly important that primary recipients instigate measures to stop the recording process if they do not want important information to be over-written. Subsequent paragraphs reviewed the legal issues surrounding the disclosure of evidence in the aftermath of an incident. This area is of particular concern when the anonymity of potential contributors might be jeopardised by the subsequent release of automated recordings.

Primary recipients are also often involved in collecting evidence from eye-witnesses. A number of techniques were therefore presented to help in this task. Different interview formats were considered. These included one to one interviews, many to one interview panels, one to many team-based interviews and many to many group discussions. Interview structures were also discussed. These included flexible interview techniques, more formal interview structures, semi-structured interviews, prompted interviews and closed response techniques. However, the information that is provided by these approaches can be subject to a number of biases that affect eye-witness testimonies. These biases stem from both cognitive factors, including post-event reconstruction, as well as the more obvious social pressures to conform to a 'group-view' of an adverse occurrence. Later sections also went on to consider ways in which more fundamental, institutional and organisational factors can influence the entire elicitation or interview process. This analysis drew heavily upon recent reports into the biases that affect the ways in which police agencies have taken and analysed eye-witness testimonies.

The closing sections of this chapter have reviewed the primary recipient's role in drafting a preliminary report. This document, typically, provides a summary of the initial data gathering tasks and may also describe the initial actions that were taken to safeguard the system. Our analysis has focussed on the way in which time limits are usually established for the presentation of these reports so that other organisations can be warned about potential failures. However, there is a natural reluctance to present what might be premature reports about commercially sensitive failures. As a result, regulatory intervention is typically required to ensure that other organisations are alerted of a potential hazard. Other industries rely upon a less rigorous approach in which the publication of safety information can be filtered or postponed until the results of a more complete investigation are compiled. The next chapter looks at the next stage in such a detailed investigation. In particular, it focuses on the reconstruction techniques that can be used to form a coherent account from the individual events that are identified in a preliminary report.

Chapter 7

Secondary Investigation

This chapter looks at the immediate follow-up to a preliminary report. It begins by examining the role of specialist incident investigators who may be called in to supplement the work of the primary recipient. In particular, it looks as the way in which they must define the scope of any inquiry. Subsequent sections describe ways in which further evidence is gathered about an incident. This is then used in Chapters 8 and 9 to reconstruct the events that contributed to an adverse occurrence.

In some cases, it may be decided that the investigation of an incident should be terminated after the publication of a preliminary report. Such a decision could be based on a preliminary risk assessment; the apparent criticality of the incident does not justify the expenditure involved in additional investigatory resources. Alternatively, such a decision could be based on the workload that must be supported by investigation teams. An incident that might receive further consideration under 'normal' circumstances might be neglected through pressure of work with other adverse occurrences. As a result, it is important to document the reasons why an investigation is stopped:

> "The reporting officer will ordinarily decide whether or not an incident is accountable or reportable. This decision cannot be an arbitrary one, but must be based on a thorough review of all evidence, as opposed to speculation, related to the incident in question and be in accordance with the requirements of the accident reports statute and the guidelines provided in this Guide. If you are certain that a particular situation is outside the scope of the reporting requirements, then the basis on which this determination was made must be thoroughly documented before the case may be omitted from the monthly submission. If there is any uncertainty as to whether or not to report an incident, it is recommended that a report be made." [233]

Clearly the decision to terminate an investigation must be monitored to ensure that it does not jeopardise an important 'learning opportunity'. Typically, the documentation that justifies such a decision should be forwarded to regional or national safety management groups for further analysis [423]. For instance, certain units might consistently assign relatively low risk levels to incidents that have the potential to cause more serious failures if any available protection measures are compromised [701].

The secondary investigation takes place after the primary recipient of an incident report has drafted their preliminary account. This document is based on initial witness statements and a cursory examination of any physical evidence. However, it is unlikely to be complete. The timelimits, mentioned in Chapter 6, that govern the production of these reports typically imply that these initial accounts will be based on partial evidence. For instance, it can take some time to extract information from automated logging systems. Similarly, lab-tests on metallurgical failures can take weeks or months to complete. It is, therefore, important that procedures are specified to coordinate any subsequent investigations.

The simplest approach to any secondary investigation is to allow the primary recipient to continue with an investigation. This has numerous potential benefits. In particular, it is likely that they will understand the local context in which an incident occurred. This is important because it can be difficult for external investigators to quickly come to terms with this situation. The primary recipient is also likely to be a trusted individual. For instance, we have described how they are often 'local champions' for the reporting system. The primary recipient may also have been nominated to perform this role by their per group. There are, however, a number of potential problems with this approach.

Some of these problems inspired the writing of this book. Later sections of this chapter will provide primary recipients with a number of techniques that can be used to support the secondary investigation of adverse occurrences. Subsequent chapters introduce analytical techniques that support more detailed causal analyses. Such written material can be supported by training courses that provide primary recipients with this information in a manner that can be tailored to the particular needs of their organisation. For example, Section 4.7 of the US Lawrence Livermore Laboratory's Health and Safety Manual states that "Training course EM2010 (Occurrence Reporting) is required for managers, supervisors, and others involved in occurrence reporting activities" [478]. The problem with this approach is that it can be extremely expensive to sustain an in-house training capability in incident investigation. In particular, it can be difficult to ensure that such guidance continues to conform with national and international guidelines. As a result, a number of organisations offer training courses that are intended to provide staff with the information and skills that are required during a secondary investigation. For instance, the American Institute of Chemical Engineers offers a course on Investigating Process Safety Incidents:

> "*You should attend if:* You are a technical professional who wants to be a team leader in critical process safety investigation situations. Process engineers, superintendents, managers, operating supervision personnel and Process Safety Management program coordinators have all found this course to be a valuable resource for developing a solid system for investigations.
>
> *You can expect to:* Learn how to be an effective process incident investigation team member or team leader. Focus on the structure and function of the investigation management system but not on root cause analysis techniques. Discover how to create a turn-key investigative management system tailored to your organisation's needs. Gain a comprehensive view spanning the scope of the investigation management system ranging from pre-planning to report generation...from structure to function. Broaden your knowledge of process related incidents, specifically vs. personnel injuries. Apply practical techniques based on up-to-the-minute reports like the AIChE/CCPS Incident Investigation Guidelines. Utilise key principles and practical skills in two stimulating workshops intended to reinforce your knowledge.
>
> *Content you can count on:* Multiple Root Cause Concept and Investigation
>
> *Methods:* Management System Development, Evidence Gathering and Analysing, Witness Interviewing, Determining Root Causes, Forming and Evaluating Recommendations, Preparing Written Reports.
>
> *Interactive Workshops:* Witness interview; Incident investigation." [24]

The use of such professional courses helps to ensure that staff are kept up to date with regulatory reporting requirements without incurring the costs associated with maintaining local training provision. However, professional courses can also incur significant costs for organisations that want to train primary recipients in secondary investigation techniques. As a result, many organisations only send regional or national investigators to these training sessions. There are also a number of further pragmatic issues that limit the use of local, primary recipients during subsequent stages of incident investigation. As mentioned, it is likely that the employee representatives, supervisors and local managers who initially receive incident reports will have a good grasp of the environmental and contextual factors that contribute to adverse occurrences. However, they may lack an awareness of regional or national safety priorities. In particular, it can be difficult for these individuals to find out about whether or not a particular incident forms part of a wider pattern of similar failures. There are also concerns that the standard of skill, training and commitment to secondary investigations will not be consistent across all branches of an organisation. As a result of these concerns, many organisations allocate subsequent investigation tasks to a small group of regional or national investigators. For example, the US Army specifies a number of detailed training requirements that must be satisfied by the relatively small number of individuals who can lead aviation incident and accident investigations [803]. These fall into two phases. In the first phase, they must complete the Aviation Safety Officer Course, they may additionally have to sit a Chemical Accident Investigator Class. They must complete classes on Blood-borne Pathogen and Hazardous Materials and Human Factors in Accident Investigations. They must also display a knowledge of the relevant military investigation guidance; AR 385-95, AR 385-40, AR 385-10, DA Pam 385-40 and USASC Investigations Handbook. Finally, they must have participated in an aircraft accident investigation orientation. The second phase of training for these 'professional' investigators includes courses on Aircraft Accident Investigation, Rotorcraft Accident Investigation, Basic Crash Survivability Investigation and on Advanced Crash Survivability Investigation. Investigators must also demonstrate proficiency in investigative and briefing skills to a board of peers and group commanders. Additionally, these individuals must ensure that they maintain 'accident investigation currency'. if more than six months passes between investigations then officers can be required to participate in a subsequent accident investigation orientation. The Chief, Aviation Systems and Accident Investigation must ensure that investigators continue to satisfy these various requirements.

As can be seen from the previous paragraph, 'professional' investigators will be better trained in incident investigation techniques than the primary recipients who initially pass on information about an adverse occurrence. However, as noted, there is a danger that national and regional inspectors will lack important local knowledge. There is also a danger that, over time, they may become isolated from the practical experience of performing the functions whose failure they must subsequently investigate. For example, many European air traffic service providers require that incident investigators have a minimum of ten years active service as controllers. However, their appointment as investigators necessarily places demands on their time that can prevent them from acting as controllers. After a relatively short period of time they must be re-trained not simply in investigation techniques but in the revised procedures and new systems that their colleagues must exploit to support their everyday tasks.

It is important to emphasise that the use of highly trained and well motivate personnel will not guarantee the overall success of an investigation. Even though a secondary investigation is performed in a reliable and consistent manner, it is still possible for the recommendations not to be acted upon. For example, incident investigators who work for the Train Operating Companies (TOC) on UK railways must satisfy the following regulatory requirement:

1. preservation and collection of evidence, including securing the scene of an accident;

2. accident investigation;

3. maintenance of confidentiality;

4. forensic and interview techniques;

5. human performance assessment; and

6. root cause analysis. [350]

However, Her Majesty's Railway Inspectorate (HMRI) enquiry into the investigation of incidents involving 'Signals Passed at Danger' (SPADs) found that "in some cases greater emphasis was placed on completing a multi-page form than getting to the root cause of the SPAD incident". This was apparent even though the actions of regional investigators were governed by railway group standard GO/RT3252 'Signals Passed at Danger':

> "Inspectors identified shortcomings in the competence of those charged with investigating SPAD incidents in some Train Operating Companies, whereas others were seeking

> to address this by suitable training in root cause analysis in order to ensure greater competency in root cause investigation techniques."[349]

The consequences of this were identified in a recent internal report that examine the failure of the HMRI to respond adequately to previous problems at the signals which were involved in the Ladbroke Grove rail crash:

> "During the almost five years preceding the Ladbroke Grove accident, there had been at least three occasions when some form of risk assessment analysis on the signaling in the Ladbroke Grove area has been suggested or proposed. The requests were: the Head of Technical Division's letter of 11 November 1996 which requested a layout risk assessment of the re-signaling (paragraph 43); the Field Inspector's letter of 16 March 1998 to Railtrack (paragraph 64); and the Railtrack Formal Inquiry of 1 July 1998 (paragraph 66). In addition there was an earlier request for details of measures taken to reduce the level of SPADs in the area around SN109 recorded in the Head of Technical Division's letter of 1st March 1995 (paragraph 39). None of these requests appear to have been pursued effectively by HMRI." [351]

These quotations illustrate systemic failures in the conduct and monitoring of the secondary investigations that are the focus for this section of the book. Local TOC inspectors were expected to investigate and report on any SPADs that were reported. However, the enquiry showed that these primary recipients had not received sufficient training to perform their duties. The HMRI inspectors were supposed to ensure that TOC inspectors investigated and acted upon those reports. However, the internal report argues that they failed to respond to the shortcomings of the TOC inspectors. In consequence, the root causes of many SPADs were not addressed before the Ladbroke Grove accident.

7.1 Gathering Evidence about Causation

Previous paragraphs have argued that the primary recipient of an incident report must be trained in investigation and analysis techniques if they are to follow-up on the information that is contained in a preliminary report. However, training courses and their supporting documentation provide no guarantees either that a secondary analysis will be performed in a rigorous and consistent manner or that any consequent recommendations will be acted upon. This chapter looks at some of the reasons why it is so difficult to build on the primary report of an incident. Subsequent chapters present a range of techniques that can be used to address these barriers to the secondary investigation of adverse occurrences

7.1.1 Framing an Investigation

One of the main decisions to be made during any subsequent investigation of an incident is to determine the scope of the analysis. This raises a number of theoretical and pragmatic problems. For example, some authors have suggested that it is possible to separate an analysis of *what* happened from the more causal investigation of *why* those events occurred [415]. This provides considerable analytical benefits. Later sections will describe how it is possible to build mathematical models of causation that link the events identified during an investigation to explain the reasons why an incident occurred. However, this division of *what* and *why* creates pragmatic difficulties for the incident investigator. For example, without some preliminary ideas about the probable causes of an adverse occurrence, it can be difficult to determine what evidence should be gathered. It is often infeasible to gather every possible item of evidence in response to a preliminary report. For example, the previous chapter cited instances in which investigators choose not to gather CVR data, because they believed that it would not make any contribution to the overall understanding of the incident. Initial informal ideas of causation, therefore, seem to play a critical role in guiding the initial investigatory process. Several of the mathematical techniques that support the causal analysis of accidents and incidents now recognise the importance of this iterative process [469]. Evidence that is obtained about the course of an incident can force investigators to revise their initial ideas about the causation of an adverse occurrence. This iterative loop is illustrated between the various phases of data gathering, reconstruction and analysis in Figure 5.1.

This chapter focuses on reconstruction techniques that help investigators to determine what happened. In consequence, we postpone an analysis of tools that can be used to support the causal modelling of *why* an event occurred until Chapter 10. As notes in the previous paragraph, however, it is difficult in practice to separate the causal analysis from the process of gathering evidence. As a result the remainder of this section introduces the notion of root cause analysis that motivates many of the techniques that we shall introduce in subsequent chapters. This is justified by the observation that any secondary investigation must uncover sufficient evidence to identify there root causes.

The US Department of Energy argues that investigations must help line management to avoid future failures by identifying the causal factors of previous incidents [207]. This implies that any investigation must detect and remove any local factors that, if corrected, might help to prevent a future failure. Investigatory boards must also identifying and describing any failures in management systems and oversight processes that allow hazards to exist:

> "Modern accident investigation theory indicates that generally the root causes of accidents are found in management system failures, not in the most directly related causal factor(s) in terms of time, location, and place. Generally, the higher the level in the management and oversight chain at which a root cause is found, the broader the scope of the activities that the root cause can affect. Because these higher-level root causes, if not corrected, have the largest potential to cause other accidents, it is incumbent on a board to ensure that the investigation is not ended until the root causes are identified. If a board cannot identify root causes, this should be stated clearly in the investigation report, along with an explanation." [207]

A key question that emerges from this analysis is; what exactly is a root cause and how does it differ from other contributory causes? Not only is this subject of considerable practical significance for incident investigators but it has also been the focus of philosophical debate for many years. Brevity prevents a thorough explanation of the various positions within this debate but it is worth reviewing two different stand-points because they have been used to guide a number of different incident and accident investigation techniques. The first of these philosophical approaches to causation was initially stated by Hume [376] and then developed by David Lewis [490, 491]. It has recently been developed into a causal reasoning tool by Peter Ladkin [470, 469], see Chapter 10, and hence is introduced in this section. Hume's contribution can be summarised in the following two definitions where objects can refer both to events and to states in a system:

> "...we may define cause to be an object, followed by another, and where all the objects similar to the first are followed by objects similar to the second. Or in other words where, if the first object had not been there, the second never had existed."

These definitions characterise what has become known as *counterfactual* reasoning. The general form of this argument is that if some event had not occurred as it did then the accident would never have occurred. This provides useful leverage because incident investigations must identify those events that we can eliminate to prevent future accidents from occurring. Lewis' contribution was to provide a mathematical model to support counterfactual reasoning; this model lies at the heart of Ladkin's Why-Because Analysis (WBA) that will be discussed in subsequent chapters. Lewis argues that *necessary causal factors* can be distinguished using a particular form of counterfactual argument. If A and B are states or events, then A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either. Lewis builds on this to consider alternative scenarios in which A did not occur and neither did B. In mathematical terms, he exploits a Kripke structure to define a nearness relationship between possible worlds. This enables us to reason about the nearest possible world in which A and hence B did not occur. All of this would be of only academic interest if it were not for the strong parallels between Lewis' philosophical approach and the activities of secondary incident investigation. For instance, there is a very real sense in which investigators are look for the closest possible world, i.e. the minimal system change, that would prevent an accident from being caused. The importance of the counterfactual approach is also illustrated by the US Air Force's definition of a causal factor:

> "A cause is a deficiency the correction, elimination, or avoidance of which would likely have prevented or mitigated the mishap damage or significant injury." [794]

The second line of theoretical thought on causation stems from Mackie's work on singular causality [508]. This is included within our analysis because Johannes Petersen has recently extended Mackie's work to analyse the ways in which operators respond to incidents and accidents [677]. Mackie argues that a cause (in the singular) is a non-redundant factor which forms part of a more elaborate *causal complex*. It is the conjunction of singular causes within the causal complex that leads to a particular outcome. Crucially, the causal complex is sufficient for the result to occur but it is not necessary. There can be other causal complexes. If any of the necessary causal factors within a causal complex are not present then the effect will not be produced. However, Mackie argues that it is a subjective decision by the investigator if they attempt to identify a single cause within the collection of necessary causes of a causal complex. He goes on to develop the notion of a causal field that describes the normal state of affairs prior to any incident. Investigators try to identify the causes of an incident by looking for disturbances or anomalies within the causal field. This causal field is, therefore, a subjective frame of reference that individuals use when trying to explain what has happened in a particular situation. If a cause does not manifest itself within the causal field then its influence is unlikely to be detected. This is important because Russell points to the uncertainty of any causal analysis that is based on partial observations of 'causal' sequences [716]. He argues that if we see a stone beside a broken window then we can never be absolutely sure that the stone caused the window to break; a bird may have flown into the glass or there may have been an inherent weakness in the material etc. Mackie's work explains this by suggesting that an individual's interpretation of cause depends upon the subjective frame of reference determined by their causal field. As before, this analysis would not be particularly significant if it had not been used to guide the causal analysis of incidents and accidents. For instance, Petersen's work builds on previous studies by Rasmussen [695] and Lind [493] in which they advocated that any analysis of system failure must be grounded on the functional structure of the system because this provides what Mackie describes as the causal field. The notion of a causal field also has strong implications for our previous discussion about the iterative nature of evidence gathering and causal analysis. For example, if an investigator develops an initial view about the causes of an incident then they may restrict their view of the causal field only to those system behaviours that provide evidence about those causes. Several of Mackie's ideas are reflected in the UK Health and Safety Executive's guidance on the incident and accident analysis that support railway safety cases:

> "There is much evidence that major accidents are seldom caused by the single direct action (or failure to act) by an individual. There may be many contributing factors that may not be geographically or managerially close to the accident or incident. There might also be environmental factors arising from or giving rise to physical or work-induced pressures. There is often evidence during an investigation that some of the contributory factors have been observed before in events that have been less serious. Accident and incident investigation procedures need to be sufficiently thorough and comprehensive to ensure that the deep-rooted underlying causes are clearly identified and that actions to rectify problems are carried through effectively. For such arrangements to be adequate under the Regulations, it is essential that incidents that have a potential to endanger people are examined effectively and that those that could lead to more serious consequences are treated with similar rigour to accidents that actually do cause harm." [350]

The idea that accidents and incidents are not caused by single factors has strong parallels with Mackie's causal complexes. The argument that accident and incident investigation must be thorough and comprehensive enough to identify possible causal factors also reflects Mackie's work on causal fields.

There are strong similarities between the work of Lewis and Mackie. However, Lewis' work focuses on more narrowly on necessary causal factors while Mackie's work on causal complexes focuses on conditions that are sufficient for an outcome but which are not necessary. The same effect may be achieved by several other causal complexes. This difference has profound practical implications. Lewis suggests that it is possible to avoid incidents by blocking the necessary and sufficient causes of failure. Mackie suggests that the best that we can do is to expand the scope of our causal field to provide a better view of a causal complex. There can, however, be little assurance that the same incident will not recur in ways that we have not been able to predict from our examination of a single causal complex. This debate, therefore, has strong similarities with the different positions adopted by Sagan and Perrow. As we have seen in Chapter 1, Perrow's work on normal accidents suggests that it is impossible to entirely engineer out certain forms of failure that are inherent in complex, tightly coupled systems [675]. In contrast, Sagan's initial position has been to argue that high-reliability organisations can systematically address the causes of failure in complex, technological systems [718].

Part of the motivation for introducing this theoretical material has been to try to clarify the underlying distinctions that often become lost in the plethora of competing definitions that have been proposed for everyday terms such as 'root cause' or 'contributory factor'. It is also important to stress that many incident investigators introduce further distinctions that build on, or arguably confuse, the concepts introduced by Lewis and Mackie. For example, the US Department of Energy introduces the concept of direct causes [207]. A direct cause is the immediate events or conditions that led to the accident. An example might be the contact between the chisel bit of the air-powered jackhammer and the 13.2 kV energised electrical cable in a sump pit that is being excavated. The US Department of Energy argues that "while it may not be necessary to identify the direct cause in order to complete the causal factors analysis, the direct cause should be identified when it facilitates understanding why the accident occurred or when it is useful in developing lessons learned from the accident" [207]. This notion of directness is a recurrent theme in many investigatory handbooks and manuals. It is also often referred to in the distinction between proximal and distal causes [486]. However, it can be difficult to explain this notion of directness in terms of the models developed by Lewis and Mackie. In some respects these proximal, direct causes are both necessary and sufficient. Using a counterfactual argument, if the chisel bit had not hit the 13.2 kV energised electrical cable then the accident would not have occurred. However, a counterfactual argument at this level provides few insights for the secondary analysis of an adverse occurrence. It can also be argued that from the point of view of the outcome, a direct cause is sufficient but not necessary. There may be a number of other direct causes for the resulting shock that was delivered to the jackhammer operator. In order to account for such paradoxes, several authors have introduced a distinction between general and singular causality [677]. Singular causality refers to relations between the particular set of events that were observed during an incident. In the case of the Department of Energy example, the direct causes were singularly necessary and sufficient for that particular adverse occurrence. General causality refers to relations between more abstract types of events that could lead to several different instances of the same failure. In the previous example, the direct causes of the particular failure were sufficient for the incident to occur but were not necessary in terms of the general outcome. There may have been several different ways in which the accident could have occurred. Clearly, incident investigation must focus on general causality if it is to prevent the outcome from recurring. Experience shows, however, that many accidents occur because safety managers focussed on the singular causes of particular failures [701].

Secondary analysis, typically, proceeds by the iterative formation and validation of various hypotheses about the causes of an incident. This validation, in turn, depends upon gathering evidence that is then used to reconstruct the events leading to an adverse occurrence. Three classes of causal factors can be identified amongst these 'events':

- *Contextual Factors: neither necessary not sufficient.* Contextual factors are events or conditions that did not directly contribute to an incident. There are many reasons why these events are considered during an incident investigation and why they can be included in the synopses that often support final reports about the causes of an incident. Firstly, they help to set the scene and establish the context in which an adverse occurrence took place. Secondly, they can help to establish that certain factors were NOT significant in the events leading to failure. For instance, incident reports often state that meteorological conditions were favourable. Adverse weather conditions might then be excluded as potential causal factors. Thirdly, although contextual factors may not have contributed to the particular view of an incident, they may play a more active role within a general analysis of alternative causes of an incident. For example, the fact that a platform was wet may not have contributed to a particular fall, however, it remains a potential cause of future slips.

- *Contributory Factors: necessary but not sufficient.* Contributory factors are events or conditions that collectively increase the likelihood of an accident but that individually would not lead to an adverse occurrence. These are the 'banal factors' in Reason's observation that "... a detailed examination of the causes of these accidents reveals the insidious concatenation of often relatively banal factors, hardly significant in themselves, but devastating in their combination" [700]. Contributing causes can be thought of as latent conditions that, alone, are insufficient to cause a failure but which were necessary for it to occur. For example, disabling a necessary protection mechanism can create the potential for a triggering event to have more serious consequences. Similarly, the failure to erect barriers or to post warning signs can contribute to an adverse occurrence. It is important not to underestimate the importance of these contributory factors as they often have the greatest general significance for future failures. It may be difficult or impossible to predict all of the catalytic events that can lead to a failure. However, the consequences can be reduced by ensuring that contributory factors are adequately dealt with in the aftermath of an incident.

- *Root Cause: necessary and sufficient.* Root causes capture Lewis' notion of causation established by counterfactual reasoning. If a root cause had not occurred in the singular causes of an incident then the incident would not have occurred. If a root causes were corrected then that the same incident would not recurr. However, as noted above, we can also introduce the stronger notion of a general root cause. These are causes that represent the globally necessary and sufficient causes that go beyond the immediate direct causes of an incident, as defined by the US Department of Energy. It is important also to emphasise that root causes can be formed from several contributing causes. This captures part of Mackie's vision of a causal complex. They are higher-order, fundamental causal factors that address classes of deficiencies, rather than single problems or faults. For example, the HSE stress that:

  > "In these criteria the term 'root causes' includes consideration of management's real and perceived messages to workers, environmental and human factors, as well as plant failures and inadequate procedures. Human errors arising from poor operating conditions, procedures, management expectations or plant design are not root causes; the predisposing factors are." [350]

  The root causes of an incident might, therefore, include the failure to implement a safety management system. Individual contributory causes might then involve failures to: define clear roles and responsibilities for safety; ensure that staff are competent to perform their responsibilities; ensure that resource use is balanced to meet critical mission and safety goals; ensure that safety standards and requirements are known and applied to work activities; ensure that hazard controls are tailored to the work being performed; ensure that work is properly reviewed and authorised [207].

It is important to notice that the 'exacerbating factors' introduced in Chapter 9 does not fit naturally within these distinctions. Having raised this caveat, it is important to note the significance of the final point in this list, which focusses on managerial root causes. Safety-critical systems are, typically, designed with defences that are based upon the premise of causal independence. In order for an accident to occur any technical failure or human error involving a production systems would have to 'circumvent' the available automated protection systems. It would also have to breach the numerous physical barriers that are usually erected to protect personnel and equipment. However, the managerial root causes of many incidents often conspire to overcome these 'independent' defences. As Reason notes, the Bhopal disaster showed that "three supposedly independent defences failed simultaneously: a flare tower to burn off the deadly methocyanate gas; a scrubber to clean air emissions and a water sprinkler system to neutralise the remaining fumes" [701].

7.1.2 Commissioning Expert Witnesses

The previous sections has argued that the scope of a secondary investigation is defined by an iterative process in which investigators form hypotheses about the root causes of an incident. These hypotheses are then validated by gathering relevant evidence. However, this evidence may reveal inconsistencies that force the investigator to revise their initial hypotheses. They must then, in turn, seek further evidence to validate their new ideas about the causes of an incident.

Chapter 6 introduced some of the problems of evidence gathering that effect the initial investigation of an adverse occurrence. Many of these problems, such as the difficulty of gathering and interpreting eye witness statements, also affect subsequent enquiries by trained investigators. Other problems stem from the iterative process of formatting and validating hypotheses. For instance, previous chapters have introduced the *confirmation bias* that makes individuals more likely to accept evidence that supports a hypothesis. It also makes them more likely to ignore evidence that is inconsistent with their initial views. Other forms of bias effect the secondary analysis of incidents by individuals working within the organisations that are under investigation. For example, attribution errors occur because individuals are more likely to attribute the causes of failure to situational aspects if they are potentially implicated in that failure. However, if they are not themselves implicated then they are more likely to look for evidence that others were to blame rather than look for wider contextual factors [121]. It is difficult to avoid what are often implicit biases. Investigators may not be aware that such factors influence their behaviour. These biases are often exacerbated by their omission from many of the courses that are intended to train incident investigators.

Further problems complicate the secondary investigation of adverse occurrences. For example, the primary recipient of an incident report usually does not have time to call for specialist reports in the immediate aftermath of an adverse occurrence. However, expert witnesses are often solicited after a preliminary report has been published. In most cases, these witnesses help to mitigate the biases mentioned above. Attribution errors can be addressed because expert witnesses may take an independent view of the investigatory agencies role in any incident. Confirmation biases are resolved because expert witnesses can use their experience to look beyond the initial hypotheses being proposed by incident investigators. Of course, a more cynical view is that these sources of additional evidence may do little more than to bolster or confirm these preliminary judgements about the causes of an incident [410]. Such cynicism contrasts sharply with the guidelines that determine the role of expert witnesses within boards of inquiry:

> "Expert witnesses also may be called to testify on selected topics to assist the Chemical Safety and Hazard Investigation Board in its investigation. The testimony is intended to expand a public record and to assure the public that a complete, open objective investigation is being conducted. The witnesses who are called to testify have been selected because of their ability to provide the best available information on the issues related to the chemical incident, or who had direct knowledge of the events leading up to the incident." [161]

This quotation stresses the positive role of expert witnesses in helping to determine the causes of incidents and accidents. However, things are not always so clear cut. For instance, the evidence of one group of experts can often be rebutted by evidence from their colleagues. As a result, regulatory organisations often publish explicit advice about the role of scientific evidence in safety assessments and risk analysis. For example, the UK Health and Safety Executive have considered this issue in a number of recent studies [330]. Although the remit of these enquiried has extended well beyond the scope of secondary incident investigation, some of the interim findings are applicable within this context:

> "Identify trusted, independent parties who your audience are likely to turn to for advice, or from whom they will form their opinions. Get them on board early. Conflict among experts will always damage credibility." [328]

If such advice is not heeded then the consequences for any incident investigation can be profound. Any subsequent litigation will be reduced to a dispute between the relative credibility of expert witnesses. In such circumstances, courts rule on the weight of the scientific and technical evidence that is presented to them. They must assess the credibility of expert evidence.

> "The administrative law judge who heard the case decided that Dr. Hochman's 'opinion is entitled to considerable weight'; nevertheless, he further decided that the opinion testimony of the Secretary's three experts about breaks before the rope snapped is of 'greater value'." [643]

If outside opinion is to be relied upon, it is therefore essential that incident investigators seek advice from a well qualified source. Most expert witnesses gain a reputation for their work within a particular area of technical expertise. As a result, information about their particular skills is often exchanged by incident investigators who need particular services. However, if investigators cannot find an expert witness through the recommendations of their peers then there are a number of alternative techniques that can be used. For example, some experts advertise their services in trade publications or directly through promotional flyers that are sent to lawyers and investigators. There is an obvious concern to validate the credentials of such individuals. One of the main means of achieving this is to consult the national professional association of the discipline concerned. However, this does not provide a guarantee of competence.

It is important to emphasise that expert witnesses do not simply need to be skilled within their own domain. There is an obvious requirement that they have some understanding of the legal framework that supports their role within a safety system. This is not always as straightforward as it might appear. For example, some aspects of the English and Welsh legal system can appear 'surprising' to potential witnesses:

> "The Royal Commission on Criminal Justice was concerned that because of the rules on hearsay evidence, an expert witness may not, strictly speaking, be permitted to give an opinion in court based on scientific tests run by assistants unless all those assistants are called upon to give supporting evidence in court. It seems to us that this rule is badly in need of change. The Law Commission agrees, and recommends that the prosecution and the defence should give advance notice of the names of anyone who has supplied information on which an expert will rely, and the nature of that information. The expert could then base any opinion or inference on the information supplied by any such person, without the party having to call that person, unless the court directs otherwise on application by any other party to the proceedings. This should result in a reduction in pointless cross-examination of experts' assistants." [477]

It may seem paradoxical to stress this issue again, when many of the reporting systems that we have considered are both voluntary and non-punitive. However, expert witnesses are most typically called to support the analysis of incidents within 'proportionate blame' systems. It is also important to remember that even 'no blame' systems must operate within the law.

Expert witnesses must possess a range of further skills in addition to their domain expertise. They must also be able to explain their insights to the many different groups who can have a stake in the results of an incident investigation. Spohrer and Maciejewski [754] illustrate this point when they stipulate ten commandments for expert witnesses. They focus on the role of expert chemists during litigation. However, their advice is generic. The following guidelines re-interpret them for incident reporting:

1. *Know the proper standard for admissibility of your testimony.* In certain areas, there are standard tests for establishing particular hypotheses. For example, in the United States post-incident examinations of trucks and buses are, typically, based around the Commercial Vehicle Safety Alliance (CVSA) criteria. There are also Office of Motor Carrier (OMC) inspection guidelines that must be followed by any independent expert. These guidelines provide a set of standards that can be used to determine whether, for example, the brakes on a commercial vehicle were satisfactorily maintained before an incident. In other areas, things are less clear cut. For example, there are several competing theories about the impact of workload on human decision making [863]. As a result, experts may use one of several approaches to determine whether or not this was a factor in a particular incident.

2. *Do your homework.* There is a legal trick which goes as follows. The lawyer asks the expert witness if they agree that some standard text is an authoritative source on a particular topic. If the expert witness agrees then the lawyer takes them carefully through the paragraphs that rebut their evidence. The standard response to this ploy is to claim that no published source can ever be authoritative because by the time that they are published there will usually be more advanced research that could not be included. In consequence, it is important for expert witnesses to keep up to date with recent developments. For instance, the tests mentioned in the previous bullet point have recently been reviewed by the NTSB following a number of incidents in which vehicles brakes failed even though they were OMC certified [607]. Any evidence that is based on the OMC certification would have to be re-interpreted in the light of this NTSB study.

3. *Always maintain your "cool" during a deposition and at trial.* This commandment relates narrowly to the role of the expert witness within litigation and so it more difficulty to apply more generally in the secondary analysis of safety-critical incidents. However, Spohrer and Maciejewski introduce a number of important pitfalls that witnesses should be aware of [754]. For instance, they warn against the negative effects of cross-examinations that include questions such as "Have you stopped beating your wife?". These 'no-win' questions could be phrased as "Dr. Engineer, is your company still manufacturing these defective widgets?" or "Doctor, are you still performing this discredited surgical procedure?". Spohrer and Maciejewski procedure recommend that experts should never give a 'yes' or 'no' answer to such questions but should use the opportunity to restate their opinion. For instance, "I disagree with your assumption. Our widgets are among the safest in the marketplace and have been used by millions of customers without an incident..."

4. *Be an expert, not a "hired gun".* The Chemical Safety and Hazard Investigation Board's terms of reference for expert witnesses, cited above, emphasised that they are intended to convince the public that an investigation is 'complete', 'open' and 'objective' [161]. In other words, they must not simply support the existing hypotheses proposed by an investigator. Fortunately, it is relatively common to find experts who are willing to act in this independent manner. For instance, the following excerpt comes from a US Coast Guard judgement in which even the government's expert witnesses agreed with the appellant:

   > "The only testimony to be found in the record on this issue is favorable to Appellant. The sole expert witness to testify stated that he approved of Appellant's decision (Tr. 194-195). The Marine Superintendent for Ecological Shipping Corp., called by the government, testified on cross examination that he thought Appellant had made the right choice (Tr. 124-126)." [825]

   It is also important to emphasise the any initial discussions between an expert witness and an incident investigator must consider the ways in which they are to be paid for their work. It is clearly unethical to make such payments contingent on the outcome of any analysis.

5. *Request a thorough briefing.* Incident investigators must provide expert witnesses with information about the general scope of an investigation. In particular, they must provide experts with access to any necessary data. It is also important that investigators explain their reasons for engaging the services of an expert witness. The expert, in turn, must determine whether they are able to provide the evidence that is expected by the investigator.

6. *Know when and when not to 'blow your own horn'.* It is important for experts to provide the information that establishes their credibility. For instance, the following quotation comes from an NTSB investigation into a non-fatal aviation incident. Although the information seems very plausible, it is impossible to know the basis of this analysis from the report alone:

   > "According to an expert on the Long-EZ, following a loss of engine power, you must maintain flying airspeed just like a regular airplane, otherwise the canard will stall. When the canard stalls the aircraft's nose will drop 10 to 30 degrees. After the

> canard stalls, if the control stick is kept fully aft and flying airspeed is regained, the nose of the aircraft will rise." [595]

It would have been far better to state the level of expertise that backs such an assessment. This does not necessarily imply that every expert ought to be named even in minor incident reports, although this is good practice. In this case, it might have been sufficient to state the number of hours that the expert had completed on this type of aircraft.

7. *Don't guess or go out on a limb.* It is important for expert witnesses to remember that some questions defy simplistic answers. In particular, many investigations rely upon evidence derived from tests that do not provide definitive answers. The majority of scientific test provide results that are based on confidence intervals. This is illustrated by the US Occupational Safety and Health Administration's (OSHA) use of expert witnesses in assessing the risks of exposure to 1,3-Butadiene (BD). The witnesses provided the following analysis:

   > "In the Downs study (Ex. 34-4, Vol. III, H-2) the standardised mortality ratio (SMR) for all causes of death in the entire study cohort was low (SMR 80; p ¡ .05) when compared to national population rates. However, a statistically significant excess of deaths was observed for lymphosarcoma and reticulum cell sarcoma combined (SMR 235; 95(The issue of reference population selection is discussed below in paragraph (viii).) When analysed by duration of employment, the SMR for the category of all LH neoplasms was higher in workers with less than five years employment (SMR = 167) than for those with more than five years employment (SMR = 127). However, neither of these findings was statistically significant." [650]

   As can be seen, the experts are carefully to note both the problems of determining a reference population for their epidemiological study. They also state which of their findings were statistically significant and which were not.

8. *Don't talk down to the investigator or other colleagues in the investigation.* It is important to note the language that was used by the experts that are cited above. This excerpt assumes that the readers can correctly interpret the use of statistics and will be aware of some of the control issues involved in such a study. The tone is of one scientist or engineer talking to another. Although the previous citation does not show them, it also included numerous footnotes so that additional details could be obtained if the reader failed to understand some of the points that were being made. However, such references in turn assume a certain technical background and scientific expertise. This is completely appropriate given the nature of the report. However, considerable additional care is required when expert witnesses must communicate their findings to groups without more diverse backgrounds. Chapter 14 will introduce a range of techniques that are intended to address these communications issues.

9. *Don't try to be an expert on everything.* It is important that expert witnesses know the limits or bounds of their expertise. Investigators are, typically, aware of the limitations in their own expertise. This often a primary reason for the use of expert witnesses. It is also illustrated by the way in which many investigatory agencies deliberately partition the skills that they require into a number of specialist areas. Staff develop skills in a subset of those areas. For instance, the Federal Railroad Administration employs railroad inspectors who investigate possible breaches of Federal laws, regulations, rules and standards and to conduct and report on incidents or accidents.

   > "The Inspector writes reports of findings and seeks correction of unsafe conditions and may be called upon to testify as an expert witness in civil suits. The demands of these jobs are many, requiring skill in evaluation, fact-finding, report writing; comprehension and application of technical and regulatory standards; the ability to gain the cooperation of individuals and organisations; and knowledge of methods used in installation, operation, maintenance or manufacturing of railroad equipment and systems." [236]

As a result, inspectors are groups around a number of specialisations including track inspectors; motive power and equipment inspectors; hazardous materials inspectors; operating practices inspectors. As can be seen, each of these divisions carefully defines the scope of expertise for each of these 'professional expert witnesses'. It is important that a similar degree of care is taken when recruiting free-lance expert witnesses.

10. *Never sacrifice your credibility.* This might seem like little more than common sense. However, it is instructive to spend a little time reviewing the way in which a court treats expert testimony during subsequent litigation about the course of an incident. For example, the following excerpt comes from the OSHA review Commission's judgement on an appeal against a decision that went in favour of the US Secretary of Labour and against the expert opinion:

    > "Keco's argument against classifying its facility as a 'blast-cleaning room' is based primarily on the opinion testimony of its expert witness, Nicholas Corbo. We conclude, however, that that testimony is entitled to little weight... In essence, therefore, Mr. Corbo concluded that Keco's facility was not a 'blast-cleaning room' because it did not have a forced-draft ventilation system. This is not, however, how the standard defines the term. The definition in section 1910.94(a)(1)(iv) says nothing about a forced-draft ventilation system. The standard's definition is controlling here. Moreover, adopting Mr. Corbo's definition would create an absurdity in the standard. Section 1910.94(a)(3)(i) sets forth a requirement that '[b]last-cleaning enclosures [including blast-cleaning rooms] shall be exhaust ventilated in such a way that a continuous inward flow of air will be maintained at all openings in the enclosure during the blasting operation'. Yet, this standard would be rendered inapplicable to the unventilated enclosures it forbids if we were to define 'blast-cleaning enclosures' as ventilated enclosures." [644]

    Such decisions illustrate the consequences when expert witnesses lose their credibility either through a failure to apply the relevant standard or through apparent contradictions within the arguments that they present.

It is possible to add a further requirement that all expert witnesses should "keep a written record of the supporting analysis that helped in forming particular conclusions". Without this information it is impossible both to assess the validity of the witnesses conclusions or to replicate their method. This is a particular problem for the human factors analyses that frequently form part of the secondary investigation of an adverse occurrence [408]. For example, the following excerpt emphasizes the problems that high workload can create for aircrews during adverse situations. Here the term is used colloquially even though there are many more technical definitions of the concept [863]

> "There can be little doubt, however, that the high workload in the cockpit contributed to the failure of the crew to notice the abnormally high reading on the No 1 engine vibration indicator that was evident for nearly four minutes after the initial vibration. It is, therefore, recommended that the CAA should review the current guidance to air traffic controllers on the subject of offering a discrete RT frequency to the commander of a public transport aircraft in an emergency situation, with a view towards assessing the merits of positively offering this important option." [8]

In contrast, the following excerpt from a far less severe incident illustrates how expert evidence can be backed-up with information about the reconstruction techniques that support particular conclusions. Here the pilot's workload was assessed in terms of direct observations about what could and what could not be seen in a similar cockpit under similar lighting conditions. Although it is possible to argue with the interpretation of 'workload' that is being used in this incident report, the documentation of supporting evidence does provide the reader with a clear interpretation of what was meant by the human factors analysis in this context:

> "Pilot workload was evaluated whilst flying an AS 355 along a low-level route at night in full moonlight conditions. One hour was spent simulating the VFR mode whilst navigating with a half-million topographical chart and stopwatch at between 1,200 and 2,000

> feet altitude. This phase also included an assessment of the ground lighting conditions in the accident area. A further 30 minutes was spent evaluating handling and navigation in the IFR mode at 3,500 feet altitude. The following observations were noted. The flight instruments were well lit, although a variety of lighting installations exist and no comparison was possible with the accident aircraft. The cabin dome lighting was too weak for easy chart reading. (The primary function of these lights is to provide back-up illumination of the flight instruments; they were not intended for use as chart reading lights). When dimmed the dome lights had a yellow tint and the yellow coloured towns on a 1:500,000 topographical chart could not be easily identified. Minor terrain features on the chart, depicted in yellow, could not be seen in flight due to the yellow tinted light. The cabin dome light eyeball could be vectored far enough forward to shine on the pilot's left knee..." [13]

The key point here is that without such supporting information about the analytical methods that scientific and technical experts use during the secondary stages of an incident investigation then it is highly likely that their findings may be questioned during the later stages of analysis. In the worst case, their results may stand until they are examined during subsequent litigation. Without necessary information about the method and scope of the expert's techniques then it is highly likely that their insights will be discredited or rebutted by the evidence of other, equally qualified, professionals.

7.1.3 Replaying Automated Logs

Chapter 6 introduced the problems that arise when attempting to safeguard the automated logs that are increasingly being used as evidence in the subsequent investigation of adverse occurrences. This section builds on the previous introduction and goes on to consider the use of these data sources to yield important insights into the causes of near miss incidents.

It is important to emphasise that the use of data recorders to support incident investigation is not a new phenomena. The maritime industry has for a long time exploited log books, navigation charts, bell and engine order logs, course recorders and hull stress meters. However, these traditional sources of information are being supplemented by more recent developments. These include propulsion and auxiliary engine computer logs, vessel traffic service systems, Rescue Coordination Center radio transmission tapes and Automatic Identification System logs [114]. As mentioned in Chapter 6, this creates logistical problems for the primary recipients of an incident report. They must safeguard these diverse information sources and coordinate their collection for later analysis. Fortunately, a range of marine voyage data recorders have been developed to collate the various measurements that can be taken on board a vessel. These systems also ensure that they are recorded and protected in one data store so that they can be retrieved for later analysis. As Brown notes, the usefulness of these systems goes beyond their role in incident investigation; " Many companies have already taken the initiative of installing Voyage Data Recorders (VDRs) not only to obtain data in the event of an accident or incident, but also to assist in managing their fleets" [114]. The following paragraphs summarise the benefits of automated logging systems. Particular emphasis is placed on their role in incident investigation, however, we also consider some of the wider benefits that these logging systems can provide.

Most automated logging systems are introduced to provide investigators with the data that is necessary during the subsequent reconstruction of adverse occurrences. These devices were initially deployed to support air accident investigations. However, they have since been installed in a wide, and ever expanding, range og safety-critical systems. For instance, tachographs are now routinely used during the investigation of road traffic incidents:

> "(Vehicles) with the electronic tachograph capability graphically show simultaneous engine and vehicle speed, and show how a vehicle was driven for a 24-hour period. This function identifies driver compliance with speed limit changes along routes. It also profiles basic driving habits. For example, if the graph shows that the vehicles speed decreased

> suddenly but the engine speed did not, the driver may have been tailgating and had to slam on the brakes to avoid an accident." [215]

Information from such sources is not simply used to analyse human and system performance immediately before an incident. Records can also be kept to determine whether or not there is evidence of similar failures over a much longer period of time. They can also be used more pro-actively. For example, tachograph records can be used to trigger US Department of Transport violation reports if drivers exceed certain operational limits [215].

As mentioned, automated logging systems have a number of uses. Not only do they record information about system performance during potential failures, many of the applications also provide live output that can be monitored. This provides potential rescuers with direct information about the events that contribute to an incident:

> "Current generation recorders now permit a watchman monitoring distress channels to instantly play back a distress call without interrupting the recording process, even as additional voice or data signals are received. Weak, unintelligible signals can be enhanced and amplified by signal processing. This allows search and rescue workers to save lives that might otherwise be lost. Tapeless magneto-optical drive systems provide immediate playback of data when there is uncertainty concerning the exact message that was received or transmitted." [222]

This illustrates how automated logging equipment can support secondary investigations long before the analysis actually begins. By providing potential eye-witnesses with important and accurate information about the the state of an application, these systems can help observers to more accurately recall the events leading to a failure. Arguably the most obvious use of automated logging systems is to validate the testimonies of people who are involved in an incident. The following excerpt cites an NTSB summary of cases in which rail recording systems were either available to validate the crew's interpretation of events or were unavailable and the subsequent investigation had to rely on witness testimony alone. Chapter 6 has summarised the many biases that complicate the task of interpreting such evidence derived from those who are involved in an incident:

> "After reviewing the information from the trains event recorders the Safety Board investigators determined that the St. Louis Southwestern Railway Company (Cotton Belt) was lax in enforcing speed restrictions. In the investigation of a 1985 head-on collision between two Amtrak trains at Astoria, Queens, New York, Safety Board investigators performed a comparative analysis of the data from the recorders. The recorded train operator activity data was compared to crewmember statements for cab signal indications and applicable wayside signal indications to develop findings in the investigation... The investigation of a 1989 derailment with the release of hazardous materials from a freight train near Freeland, Michigan was noted as being hindered by the absence of multi-event-recorder data. The Safety Boards report stated that train-handling information was derived from what the train crew stated. The paper-tape-recorded train speed was of limited usefulness since the manner in which the train was controlled was more important than its speed. Vital information, such as quantified braking, throttle manipulation, and the chronological relationship between power-to-braking and braking-to-power, was not available". [214]

It is important not to underestimate the practical difficulties that are involved in using automated logs to validate eye-witness testimonies. As the previous citation shows, these systems do not always provide the evidence that is necessary to prove or disprove key aspects of their statements. Simply recording more data does not always provide straightforward solution. Later sections will identify some of the problems that can arise in both filtering and in interpreting the mass of data that these systems can record. You may be able to determine that the operator did issue a particular command at a particular moment in time, but no logging system will currently tell you why that command was selected.

The positive side of automated logging focuses on the use of these records to encourage future improvements in operator performance and system reliability. This provides another aspect to the way in which some organisations blur the distinction between a safety-critical incident reporting system and a more general approach to quality improvement:

> "The Navy uses recording devices as training tools to improve air traffic control operations for both ship and shore-based facilities. Operators are given the opportunity to hear themselves and see the consequences of their actions in replicated scenarios. This enhances readiness by allowing total system simulation, and by providing both individual and team training. Managers and commanders can better measure readiness, identify whether proper operational procedures are being used, and evaluate the outcome of using those procedures. Recorders offer the opportunity for students to safely learn from their mistakes in an unbiased, objective mode."[222]

The same techniques of replay and simulation that are described in this citation can also be used more directly to support the secondary investigation of an incident. Showing automated logs to an operator or eye-witness can trigger recollections that might otherwise not form part of their testimony. There is, however, a danger that such an approach may evoke a form of false memory syndrome. This is particularly apparent when the automated logs are presented through sophisticated, three-dimensional simulations. For this reason, several organisations have moved to limit the use of such reconstruction techniques during some stages of incident investigation [423]. Witnesses should only be shown replays on the equipment that they actually had available to them during the incident itself. It can be argued that this is an unnecessary restriction. Further work is urgently needed to determine whether these are valid concerns during the subsequent investigation of an adverse occurrence. However, there is often a justifiable fear that automated logging systems will not primarily be used as a safety tool. Instead they will be used to monitor employee compliance with organisational objectives and performance criteria. Later sections will describe how many automated monitoring systems are deliberately developed so that they can be customised to the requirements of the companys that buy them. The previous positive comments about the use of these systems must, therefore, be balanced against their more sinister application:

> "Competent personnel love them, while incompetent personnel loathe them. What better documentation for management to have in an incident than an exact record of actions that were (or were not) taken. Multi-tiered security systems embedded in the design of todays naval recorders prevent unauthorised access to the recorded information, thus preserving the integrity of the data for use in accident investigations or analyses. Additional features prevent the overwriting of data previously recorded on another machine. Modern recorders can also be synchronised to a universal time standard such as global positioning system (e.g., Havequick time). This allows platform-unique data to be recorded and played back in synchronisation with recording systems in other locations, thereby improving time-sensitive accident investigations." [222]

The previous paragraphs describe some of the benefits that can be obtained both for the secondary analysis of an adverse occurrence and also more widely in the operation of safety-critical systems. However, all of these benefits depend upon the deployment of the monitoring equipment. Typically, in spite of the claimed commercial benefits, there systems are not widely used unless they are backed by regulatory requirements. There are notable exceptions, however as usual, these tend to be companies that already have a high reputation for their safety management systems. The problems that arise when attempting to introduce reporting systems can be illustrated by the complex negotiations within the International Maritime Organisation (IMO). The 44th session of the IMO Sub-Committee on Safety of Navigation considered the adoption of Voyage Data Recorders (VDR). Several options were considered during this meeting:

> "The proposed options include a provision limiting the new requirement for VDRs to Ro-Ro (roll on-roll off) passenger ships on international voyages. Other options, which were submitted by the United Kingdom and supported by the European community, the United States, Canada, Australia, and New Zealand, require that all new vessels built by a certain date have a VDR and that all existing vessels install a VDR during a phase-in

> period, which will be at a later date... Some countries opposed the VDR requirement for all vessels. Japan and others stated that the carriage requirement should apply only to vessels on international voyages; Panama maintained that the VDR should only be required on self-propelled vessels." [114]

Coordinating the adoption of automated monitoring equipment is simpler when a single national regulator has jurisdiction over an industry. However, regulators must still address the problems of gaining employee trust and of convincing industry that such systems are not an unnecessary burden. Even once automated logging systems are widely deployed, a host of further problems complicate their use within the secondary investigation of adverse occurrences. These problems range from design limitations through to installation issues and the difficulty of maintaining often complex digital equipment in potentially 'hostile' environments:

> "There are no (Federal Railroad Administration) requirements for records to be kept about recorder system specifications, or applicable readout software... While a readout of the data is required every 92 days for tape-based recorders only, there is no requirement (for any type of recorder) to test the sensors or other system components or to verify that accurate data is actually being recorded. Furthermore, under current FRA regulations, microprocessor based recorders are not required to be readout, tested, or examined unless the recorder itself indicates a fault from its self-diagnostic test... (These tests) detect the presence of certain sensors, they cannot test the validity of the signals coming from the sensors. If an errant axle generator continuously sends a signal representing 0 mph, the self-test feature will not detect a malfunction. Failures such as this one may never be detected, because there are no requirements to ever read out, test, or evaluate this type of recorder. Additionally, self-test features can not detect improper programming or set-up of the recording system." [214]

The following list builds on this analysis and identifies a number of more detailed barriers to the effective use of automated logs in the secondary analysis of adverse occurrences. A common thread running through each of these items is that the installation of particular devices and the protection of their data in the aftermath of an incident do not provide any guarantee that reliable information will be obtained about the causes of failure.

Automated recording devices may simply fail to operate. In some ways this simplifies the investigators task because they do not have to piece together partially corrupted data. On the other hand, they are left to determine the reasons why such critical equipment was not being operated. Chapter 6 provided a number of examples in which data recorders were either sabotaged. It also described some comparatively rare incidents in which equipment was lost as the result of extreme forces in the aftermath of an incident or accident. The failure of most data recorders, however, often stems from more complex causes:

> "During normal operation of the system, when aircraft power was applied, the tape transport would run for 1 minute without recording data to enable different flight sectors to be separated upon replay. The system would then enter standby mode with no tape motion and the mechanical indicator on the control panel indicating 'STBY'. Once the crew had started both engines, as part of the startup procedure, they would select the aircraft generators to 'ON'. This action would switch the tape transport on, initiating the recording of data and setting the control panel indicator to 'RUN'. From this point a further 2 minute period was required to allow the Built-In Test Equipment (BITE) to detect and indicate a system fault. A later item in the checklist required the crew to ensure that the control panel mechanical indicator was showing 'RUN' and that the BITE fault indication was extinguished. A fault in the track change sensing of the tape transport of G-ATMI's recorder had allowed the tape to run off the end of one reel, become stuck to the tape drive capstan and then wind backwards around the capstan until it had jammed. Following the engine starts, prior to the accident take-off, the crew had selected the generators to 'ON', thus setting the FDR system to 'RUN'. However, the CVR recording showed that the checklist item to ensure normal operation of the

> FDR system had been carried out within 1 minute of switching the generators, which did not allow sufficient time for the system BITE to detect and indicate the fault in the tape transport. The position of the control panel on the flight deck was such that neither crew member would have been able to see the fault indication without turning to look over their shoulder." [17]"

The previous quotation illustrates the care with which incident investigators must investigate the sources of such failures. Although automated logging systems do not directly contribute to the causes of an adverse occurrence, their failure jeopardises the investigators ability to accurately identify those causes.

There are many ways in which automated recording equipment can fail to provide necessary information about the course of an incident. As we have seen, the design of the equipment may not record all of the parameters that are necessary during any subsequent investigation. Such problems are being addressed by the development of an increasingly sophisticated range of digital recording devices. There are also a number of common technological problems that can affect the analysis of flight data recorders. For example, many recorders fail to deal adequately with information that is buffered in a volatile store immediately prior to any adverse occurrence:

> "The Universal Flight Data Recorder (UFDR) takes flight data into one of two internal memory stores, each holding about one second of data. When one memory store is full, the data flow is switched to the other store. While the data is being fed to this other store, the tape is rewound and the previous second of data is checked. A gap is left on the tape and the data in the first store is then written to the tape, and the first memory store emptied. This whole 'checkstroke' operation takes much less than one second to complete... Thus the UFDR tape is not running continuously. The tape first accelerates from stationary to 6 inches per second to read the previous data block, leaves an inter-record gap and then writes the new data block. The tape then slows and rewinds ready to begin the next 'checkstroke' operation. A total of 0.48 inches of tape is used to record one block of data and inter-record gap... When power is lost from the recorder, the data held in the volatile memory which has not been recorded on the tape is lost. As can be seen from the way in which data is temporarily stored on this UFDR and then recorded, this can mean that up to 1.2 seconds of data may be lost just before impact." [8]

The AAIB continue to report similar problems. For instance, the buffering of data by a digital flight data recorder led to significant problems for the investigators of a recent loss of control incident [18]. The data buffer was not crash protected and required electrical power to retain the contents. When it was replayed, it was also found that the recorder had an undetected fault which resulted in the random corruption of all parameters over the duration of the recording. The recorder's built-in test circuitry was incapable of warning the operators about the presence of this particular fault.

Secondary investigations must make the best use of data that is provided by automated monitoring equipment. However, experience with failures in this equipment varies considerably from industry to industry. The previous problems noted with flight data recorders do not seem to have been such a concern in the railway industry. For example, the following citation describes the NTSB's experience with these systems:

> "The actual recording device itself is seldom, if ever, at fault. In fact, none of the microprocessor recorders that the NTSB has had tested thus far has ever been found to have failed, be out of tolerance, or to have malfunctioned." [214]

In contrast, most problems seem to arise from the data supplied to the recording device. Anomalous or missing data often results from inoperative, incorrectly installed, or out-of-calibration sensors. Many of the NTSB's concerns about this class of recording system focus upon the quality of maintenance that these devices receive.

> "The event recorders maintenance and its location within a locomotive were addressed in the Safety Boards report of the 1996 freight train derailment near Cajon, California.

> The post-accident testing of the microprocessor type of event recorder showed that one event recorder had a broken wire in the axle generator, as a result of an improper modification, and that another was improperly programmed. In addition, the self-diagnostic indicators were insufficient to fully examine the recording status of the units. The pre-accident inspections had been inadequate." [214]

Such concerns have complex organisational and regulatory causes. It is unclear whether substandard maintenance and inspection stem from a perception that these devices are not 'essential' for the actual operation of the railroad. It could also be argued that maintenance problems also stem from the inherent complexity of the monitoring devices and the relatively fragile nature of some sensors. Alternatively, better self-test functions could provide operators with a clearer indication that equipment is not functioning as intended. Further work is urgently required to resolve some of these outstanding issues.

Current generations of automated data recorders offer great flexibility. For example, in the rail and maritime industries it is possible to configure or progam these devices to monitor and record information about events that are of specific interests to the companies that operate them. In aviation this has led to the growth of Flight Data Monitoring (FDM) and Flight Operational Quality Assurance (FOQA) programs. The growth in the scale and complexity of the devices that support these initiatives can be illustrated by the increasing number of parameters that are simultaneously recorded. The first generation systems read from 5 to 30 parameters from metal foil storage. More recent versions of what have become known as Quick Access Recorders, to distinguish them from accident recorders, now sample from 200-300 parameters [92]. It is hard to underestimate the technical challenges that these systems can pose. As mentioned, microprocessor recording systems are typically configured to meet the customers specific requirements. As a result, it is likely that one operators requirements will be different form anothers. Additional problems arise because an individual operator can change their own requirements over time. Recorder manufacturers also update and revise system configurations as new technology is introduced. Incorrect setup or programming can lead to certain parameters being recorded incorrectly or not being recorded at all [214].

Problems can still arise even if sensor signals are reliably received by a recording system and the system is correctly configured to receive those signals. In particular, a significant amount of incident data has been lost in recent years by improper or incorrect handling procedures while the data is being prepared for analysis. These handling problems take a variety of forms. For example, recording media have been placed too close to strong electro-magnetic sources. They have also been placed in direct sunlight and even accidentally immersed in water so that even relatively resilient housings have been compromised after the equipment has been removed from the system that is being monitored. Further problems have arisen during the process of transferring data from a primary recording medium to a secondary or back-up source:

> "When the copy tape was first replayed it yielded 60% bad data, making analysis of the readout difficult, and it was not possible to determine whether this data contained the landing. This copy tape was then replayed by AAIB using both the original Copy Recorder and the AAIB replay facilities, and this yielded 95% good data for the incident. Analysis showed that this data ended when the aircraft touched down, giving incident data for 116 seconds additional to that recovered directly from the Universal Flight Data Recorder (UFDR). The copying process appeared to have repositioned the tape in the UFDR incorrectly after the down load, allowing the final approach data to be overwritten by the engine ground runs." [12]

Fortunately in this incident, the primary source was uncorrupted and the analysis could proceed as planned. However, such incidents reinforce the point that simply gathering and recording data does not guarantee that it will survive in an uncorrupted form until an eventual analysis.

The increasing flexibility and capacity of the recording systems that can support incident analysis also raises further problems for the interpretation of the data that they collect. Increasingly, these problems are being addressed by a range of sophisticated reading tools that provide and visualisation capabilities:

> "The Decision Support System is a uniquely designed relational database system that allows for extraction of information such as what-if and queries of a large number of events stored in the system. FOQA II uses high fidelity visualisation and simulation whenever feasible, to display a situation or an analysis. Visualisation is 3-dimensional. The Visualisation and Simulation can be used to display and replay Allied Signal Enhanced Ground Proximity Warning events using a photo realistic terrain database." [92]

Later sections will consider the use of simulation and visualisation techniques to support incident analysis. For now it is sufficient to realise that different configurations will be required so that any reader can correctly interpret the different configurations of recording devices: As a result; "a recording system installed on a particular operators locomotive requires a readout program that is unique to that operator" [214]. If a similar recording system were to be installed on another operator's rolling stock then there is a good chance that it would require a different readout program. Some rail recorder manufacturer support more then 50 different configurations, each requiring different software to properly extract the data. If a recorder is analysed using an incorrect or outdated reader then it is likely that some of the resulting data will be corrupted.

Chapter 5 briefly outlines the benefits of automated logging systems as a means of monitoring performance and, thereby, detecting potential incidents. This chapter also described the personal, social and organisational barriers to the introduction of these devices. Chapter 6 went on to identify the problems that occur when primary recipients have to safeguard automated logs in the aftermath of an incident. They must protect systems from deliberate sabotage. They must also prevent the inadvertent damage to logs when there are considerable pressures to resume operation. On looped recording devices, they must intervene ensure that critical data is not over-written. This section has focussed on the challenges that arise once data has been retrieved in the aftermath of an incident. Technical problems in the configuration of sensors, of the recording media or of playback devices can corrupt automated logs. In particular, installation and maintenance problems can reduce the effectiveness of these devices as reliable sources of information about the causal factors behind adverse occurrences .

7.2 Gathering Evidence about Consequences

The previous section argued that these are mutual dependencies between the search for evidence and the formation of causal hypotheses. The search for evidence is often guided by hypotheses about the root causes of an incident. This evidence, in turn, helps to refine preliminary hypotheses. This is only on aspect of the situation that confronts many incident investigators. The previous definitions of contextual factors, contributory factors and root causes looked at the events which occur before an incident. However, evidence about these events can often only be obtained by looking at the events immediately after an adverse occurrence. From the previous argument, this implies that causal hypotheses are effected both by evidence about those events that *contributed* to an incident and by those events that occurred as a consequence of an incident. For example, a recent NTSB report found that metal fractures could only have been caused by a container being loaded on top of a 'foreign object' as it was installed on a railcar. There was no direct evidence of the foreign object but it was argued that such a cause is the only explanation for the consequences that were observed:

> "Investigators found that the cracks discovered in Thrall cars were not related to car age, mileage, service pattern, maintenance, or previous repairs but to stress forces caused by the presence of a foreign object on the floor of these cars. The UP inspections of Thrall cars that ultimately prompted EW-161 provide additional evidence of this phenomenon. Further, inspections of 1,653 cars still in service since EW-161 was issued, in December 1997, have resulted in the repairs of 27 Thrall double-stack container cars, all of which had damage due to foreign objects. No evidence suggests that any of the weld failures found by the FRA or during the EW-161 inspections were the result of any other condition or phenomenon. Therefore, the Safety Board concludes that a direct causal relationship exists between the misloading of a loaded container on top of a hard

> foreign object and the weld failures at the floor shear plate to bulkhead bottom angle on Thrall 125-ton deep-well double-stack cars." [612]

This quotation illustrates many of the complexities that arise during the secondary investigation of adverse occurrences. Firstly, the lack of direct evidence for the foreign object forces the investigator to form and test a number of alternative hypotheses. The report tells us that the cracks were not related to car age, mileage, service pattern, maintenance etc. Although the report does not inform us of the techniques that were used, the reader must assumed that considerable efforts were made to obtain the necessary evidence to eliminate these possible alternatives. We are then left with the hypothesis that a foreign object caused the weld failures. This illustrates another form of causal reasoning which is similar to the counterfactual approach of Lewis. The previous quotation provides an example of a more general form of argument known as 'reductio ad absurdum'. This proceeds by assuming the opposite of the thing that you want to prove. In this case, we assume that the fractures were caused by the age of the car or by mileage. The investigator then looks for evidence to show that it is impossible or irrational to believe these alternative hypotheses. For example, by showing that the car was only three years old or that it had done significantly less miles than other comparable cars. By eliminating all of the alternatives and by proving that it is incorrect to assume otherwise, you indirectly provide support for the thing that you want to establish.

7.2.1 Tracing Immediate and Long-Term Effects

The secondary investigation of the consequences of an incident is not simply intended to gather clues about the root causes of an adverse occurrence. In many cases, this information is used to assess the severity of the incident. Chapters 1 and 2 have introduced the problems associated with any estimate of the potential 'cost' of an incident. However, a qualitative estimate of the consequences of an incident can be given by some (qualitative) function of the proximity to a particular event and the losses associated with that event. The severity of an incident is most easily assessed when there are objective physical measures these values. For example, the nearness to a airspace collision can be measured in Cartesian space. The consequent loss associated with that event can be represented by the number of lives that are threatened by such a collision. These criteria were used to calculate that the following incident should be ranked as a category C air proximity violation:

> "Shortly afterwards, the Mentor heard the Air Arrivals controller announcing that he had turned SAB 603 onto 310 degrees and immediately informed him that a British Airways aircraft, callsign BAW 818, was also airborne on a 'Brookmans Park' SID. The two controllers then instructed their respective aircraft to alter heading and noted from their Air Traffic Monitor (ATM) screens that the two aircraft symbols were very close. Subsequent calculations revealed that the minimum separation was 200 feet vertically and 0.16 nm horizontally when the highest aircraft was at 2,400 feet agl. All the flight crews involved in the incident complied fully and correctly with ATC instructions. At the time of the incident, both SAB 603 and BAW 818 were in cloud and none of the crew members in either aircraft saw the other." [15]

This incident is relatively straightforward. The air traffic controllers' who contributed to the incident were almost immediately made aware of the consequences of their actions. This simplifies any secondary investigation because the individuals who are involved in an incident can help to piece together the events both before and after an adverse occurrence. This task is made far more difficult when the individuals and teams that contribute to the causes of an incident, have little or no idea about the impact of their actions. Such incidents are particularly incidious. There is a danger that the groups who contribute to an initial failure will not alter their behaviour unless they are made aware of the consequences of their actions. These sorts of failures are typified by maintenance incidents. Two frequent scenarios reappear in the incident reports that are submitted in many different industries. In the first scenario, engineers fail to correctly reassemble some sub-component that is then placed in service for a prolonged period of time. This component might fail at any time given the presence of some catalytic event. The maintenance problem is only identified during the next scheduled maintenance interval when the original engineer might have incorrectly assembled many other devices [502]. The second scenario is illustrated by the following example. In this incident, maintenance procedures are not completed. As a result, there is a system failure and an accident is only avoided by a number of fortuitous circumstance:

> "Following an indicated loss of oil quantity and subsequently oil pressure on both engines, the crew diverted to Luton Airport; both engines were shut down during the landing roll... The investigation identified the following causal factors: 1.The aircraft was presented for service following Borescope Inspections of both engines which had been signed off as complete in the Aircraft Technical Log although the HP rotor drive covers had not been refitted. 2.During the Borescope Inspections, compliance with the requirements of the Aircraft Maintenance Manual was not achieved in a number of areas, most importantly the HP rotor drive covers were not refitted and ground idle engine runs were not conducted after the inspections. 3.The Operator's Quality Assurance Department had not identified the non-procedural conduct of Borescope Inspections prevalent amongst Company engineers over a significant period of time." [12]

This separation of causes from consequences creates considerable problems for investigators. They must work backwards from the aftermath of an incident to assemble the evidence that will eventually identify and explain the root causes of failure. The following quotation provides a further example in which the causes of an incident are separated from its consequences. In this case, medical staff initially had no idea that a syringe had been filled with the wrong drug. Only 'in retrospect' were they able to test the device and piece together the causal sequence that caused the problem. This example also illustrates how such a separation also creates immediate problems for the staff who must respond to the consequences of any failure:

> "Unknown nurse prepared 'ephedrine' labelled syringe the day before and left in OB operating room for emergency use, as was the usual practice at this hospital. On day of surgery patient had hypotension after spinal, we gave 'ephedrine' syringe and had intermittent unusual responses of severe ectopy, tachydysrhythmia, hyper and hypotension. There was delayed recognition that the 'ephedrine' syringe may have been the problem because patient had some more benign ectopy and tachycardia prior to giving 'ephedrine' and after giving 'ephedrine' the response was intermittent not immediate and lasting. Post op patient had small MI but is in no way impaired and otherwise fine and baby is fine. In retrospect the syringe became suspect and was tested and found to contain epinephrine rather than ephedrine." [755]

It is important to realise the impact that such situations can have upon the individuals who are involved. The nurse may well have realised that they could be implicated in any subsequent investigation. This can create considerable personal distress. An individual sense of guilt can be exacerbated when the staff who are involved in the causes of an incident cannot help to mitigate its consequences [7]. Instead, they must rely upon the skill and knowledge of their colleagues to rectify an adverse situation. Previous chapters have emphasised the complex, systemic causes of failure. It is interesting to note, therefore, that this voluntary, anonymous incident report focuses on the actions of a single nurse. It ignores the managerial and organisation issues surrounding the preparation of a labelled syringe on the day before the procedure. These issues were, however, commented on by a number of anaesthetists who responded to the original incident report. The separation between causes and consequences also raises a number of more complex organisational issues. There can be a delay while investigators attempt to re-establish the causal chain that links the consequences of an incident to its root causes. This creates an interregnum in which organisations can suppress or destroy evidence. They can prepare a legal defence or may even take precipitous action to forestall legal action, such as sacking individual members of staff [701].

The secondary investigation of an incident must monitor and record the consequences of any adversed occurrences. These consequences help to assess the criticality of the event. They can help to identify causal factors. This, in turn, helps investigators to ensure that the individuals, systems and organisations who are involved in a failure are ultimately informed on the consequences of their interaction. However, the investigator's tasks are further exacerbated when the consequences of an incident develop over a prolonged period of time. Air proximity incidents are relatively simple; any consequent loss of separation can be measured relatively quickly after it has occurred. Other incidents are far more complex. In particular, it can be extremely difficult to predict the long term consequences of medical incidents in which quality of life must also be considered:

> "We performed continuous spinal anaesthesia for femoro-crural bypass surgery. During the operation the patient had no pain, but was still able to move her legs... Towards the end of the operation, with regard to postoperative analgesia, we wanted to give intrathecal morphine. But instead of 0,1mg as intended, an overdose of 1,0mg morphine was injected together with another 5 mg of hyperbaric bupivacaine. The error was immediately detected. SpO2 remained at 98% with 4l/min nasal O2. Naloxone 0,08mg IV was given, followed by a continuous infusion (initially 0,2mg /h, then decreased according to clinical symptoms). The patient stayed in the Post-Anaesthesia Care Unit for the next 18 hours. During this time there occurred no respiratory complications. A slight pruritus and a 12 hour amnesia, were the symptoms experienced by the patient. She was informed about the incident and satisfied with the outcome." [755]

The causes of this adverse occurrence were determined 'immediately'. However the consequences required careful monitoring for at least eighteen hours after the event. It is difficult to underemphasise the importance of such incidents for the medical community. Recent recommendations, such as those contained in the Institute of Medicine report [453], make it clear that there must be longer-term monitoring of the clinical outcomes of adverse occurrences. In particular, the point has been made that it may not be possible to predict the long term outcome on the basis of an initial post-operative assessment. Such arguments have also been expended into more general suggestions to expand the scope of clinical monitoring to increase the detection of clinical incidents. Not only must we assess the outcome of adverse occurrences on those patients that we know have suffered from inadequate care but we must also monitor the outcomes for a wider group of patients in order to improve our detection of those incidents.

This section has focussed on the geographical and temporal distances that separate the causes of some incidents from their consequences. It has been argued that this complicates the secondary investigations that must trace the complex relationships between precursors and outcomes. However, the previous examples have illustrated relatively simple cases. There are further pathological incidents in which causal events have occurred years before other organisations have suffered the consequences of failure. For example, the Watford Junction railcrash took place in August of 1996 [348]. The original signaling that was a contributory cause to the accident had been completed and commissioned between May and June 1993. Between November 1994 and the time of the accident, the HMRI made a number of attempts to arrange an inspection of the site without success [421]. The wording of a Railway Signaling Standard (SSP 20) was imprecise. This led to a speed restriction sign being placed in an inappropriate position, which gave confusing information to the train driver. This standard had been drafted and reviewed long before the accident occurred or the signaling was installed. Such a timespan creates incredible problems for secondary investigations. The companies and individuals who contributed to the design, development and maintenance of particular components may no longer be employed to support existing systems. Documentary evidence about those components may only exist in fragmentary form. As the interval between the root causes of an incident and its eventual consequences increases, there is a corresponding increase in the importance of poor safety management and weak regulation as contributory causes. These organisations, in theory, should have had ample time to detect a problem and resolve it before the incident occurred. This is not as easy as it might seem, especially if regulatory organisations are involved in the initial decisions that create the root causes of an incident. For example, the following citation described how federal authorities partly financed a signaling system that was not ultimately supported by an adequate safety case:

> "The CSX Transportation (CSXT) and Maryland Rail Commuter (MARC) had operational reasons to modify the Brunswick Line signal system: improve passenger safety

> and freight train operations by changing the method that CSXT dispatched and monitored trains, upgrade the system capacity to operate more trains with increased peak and midday service, increase the MARC labor and equipment productivity, and reduce the CSXT operating costs. Identifiable improvements, such as total trains, traincrew use, cost savings, and Centralised Traffic Control (CTC) operations, could be quantified and measured; however, the signal system modifications did not address the overall safety of the signal system for traincrew use... The Safety Board concludes that Federal funds granted for the signal modifications on the CSXT Brunswick Line to accommodate an increase in the number of MARC trains did not ensure that the safety of the public was adequately addressed. Therefore, the Safety Board believes that the Federal Railroad Administration (FRA) should require comprehensive failure modes and effects analyses, including a human factors analysis, for all signal system modifications and that the Federal Transport Administration (FTA) should revise the grant application process to require the same such analyses be provided for all federally funded transit projects that are directly related to the transport of passengers." [596]

The previous paragraphs have described how it is important for the secondary investigation not only to gather evidence about the causes of an incident but also to monitor the consequences of any failure. The outcome of an adverse occurrence provides investigators with important information about its criticality. It can also help to ensure that all of the parties who contribute to an adverse occurrence are identified and informed about its impact upon application processes. Finally, it is important to investigate the consequences of an incident because this helps to determine its criticality. There is an important caveat to this last point that we have not raised in this chapter. In particular, we will see in Chapter 10 that the risk assessments that are derived from particular incidents need not mirror the actual consequences of an adverse occurrence. For example, some organisations adopt the policy of assuming the 'worst plausible outcome' . As a result, some Air Traffic Management providers assume that if aircrews detect and resolve an air proximity violation then that incident should be treated as if the aircraft had collided because controllers failed to actively intervene to prevent a potential accident [423].

7.2.2 Detecting Mitigating Factors

The previous section has described how some of the consequences of an incident can be separated in time and place from the immediate events that lead to an incident. As a result, it can be difficult for investigators to fully assess the outcome of an adverse event until some time after it has occurred. This section investigates a number of further complications that frustrate secondary investigations. In particular, it identifies ways in which the intervention of operators and automated systems force investigators to consider alternate hypotheses about the consequences of an incident without these mitigating factors. This represents a particular extension of Lewis' counterfactual arguments [490]. We summarised his approach to causation by stating that that 'if some event had not occurred as it did then the accident would never have occurred'. Consequence analysis often takes the form of 'if some mitigating event had not occurred as it did then the accident would have been far worse'. As can be seen, therefore, mitigating actions can be though of as a form of complement to the causal actions that lead to incidents and accident.

The following incident illustrates the way in which staff and automated systems often have the opportunity to detect an adverse occurrence and intervene to mitigate its effects. If the staff had monitored the set up of the heating blanket or if they had inspected the patient's legs during the operation then the burns might have been avoided. This form of incident represents the simplest case for consequence analysis because it is difficult to see how the outcome could have plausibly been much worse given the particular heating system that was involved:

> "After surgery, burns on the foot, posterior calf, and posterior medial thigh were noted. Surgery was lengthy. Burns are second degree, requiring at this point, topical treatment. Blistered areas are 1 X 2 cm. (foot), 4 x 8 cm. (calf) and 3 x 5 cm. (thigh). Due to the size of child, he was placed on top of the blanket with the nozzle between his

> legs. The company believes the leg was too close to the nozzle, which protrudes 10 cm. into the blanket, and the hot nozzle/hot air burned the skin." ([272], Report Number 9681384-1997-00016).

As mentioned, however, the identification and analysis of the potential consequences of any incident can be complicated by the ways in which operators or safety systems intervene to mitigate the worst effects of any failure. Of course, these fortunate interventions help to avoid accidents and more serious incidents. However, they force investigators to consider a large number of hypothetical worse case scenarios in which operators and systems did not intervene to mitigate the failure. Again, there are many incidents in which this is can be relatively straightforward. For example, the worst case in the following incident is clearly that the patient could have died if the staff had not been able to offer effective cardio-pulmonary resuscitation (CPR) in time:

> "At 12:50 pm Charge Nurse entered patient's room. Patient was dusky in colour and without vital signs. Ventilator and alarms not sounding. Ventilator circuit observed to be disconnected from TRACH, ventilator producing air however pressure alarm did not sound. Circuit reconnected to TRACH, then removed to initiate manual ventilation and CPR. After circuit disconnected for CPR alarms sounded in approx 5 seconds." ([272], Report Number 221768)

The potential consequences of many incidents are, however, often less clear-cut than this example. At the extreme, an investigator might consider that an apparently minor incidents could have 'snow-balled' into a major accident involving a significant loss of life. Although this might seem to be nonsensical, it is important to remember that many major catastrophes have apparently simple root causes. The match that triggered the Kings Cross fire [247] provides an example of this. Several investigations into previous fires on the London Underground failed to understand the potentially disastrous consequences of such events. Partially as a result, safety managers focussed on putting out those fires that did occur rather than trying to eliminate the potential for a fire to start. The following incident provides a further example of this problem. The ingestion of flying insects into a vent tube forced the crew of a commercial airliner to glide towards the nearest runway. This relatively simple problem could have had disastrous consequences. The key point here is that the organisation concerned, like the London Underground, still failed to predict these consequences even though a number of similar failures had previously been reported:

> Fuel at time of departure was 56 gallons, of which 40 was in the tip tanks... Climbed to cruise altitude of 5,500 feet MSL, leveled off, turned off boost pump. Engine lost power about 1-1/2 minutes (estimated) after changing tanks... Established glide to nearest airport and commenced restart procedure...and declared emergency. Engine restarted at 500 feet AGL on short final... Landed without incident, with full power available... Cause of engine-out was determined by mechanic at FBO to be "leaf roller" (flying insect) debris packed into right tip tank vent tube, totally obstructing air flow in the vent. Tank vents...open to air at a point under the wing attachment point. There are no screens on the vent openings. The vent was cleared, and the left vent checked and also cleared of similar debris (although not completely closed), and the aircraft was returned to service..." [61]

There are many reasons why the secondary investigation of an incident report must gather evidence about mitigating events. Not only does this provide important information about potential 'worst case' scenarios using an extension of Lewis' counterfactual arguments, evidence about the defences that protect safety-critical systems. As we have seen, human operators and automated systems are often designed to provide 'defence in depth' so that if one fails to protect an application then another may successfully intervene. However, Reason argues that many incident have multiple root causes that together may combine to defeat safety measures [701]. As a result, it is imperative that we learn as much as possible both not only about those defences that succeed but also about those defences that fail to offer the intended protection during particular incidents. For example, the following incident illustrates a situation in which a warning display in the cockpit was able to back-up the human surveillance of the cabin staff. It also illustrates how fortuitous circumstances, in particular the availability of additional company personnel on-board, often help retrieve adverse situations:

> "At FL330 had momentary [warning] message 'Door Left Aft Cabin,' meaning door 2L was not fully latched. Message cleared itself, then reappeared. (Got message a total of 4 times.) Contacted purser to have her ensure no one was tampering with door. She said there was a female passenger who had been acting very strangely since leaving [airport]... Through an interpreter...passenger admitted to having attempted to open door. [Crew] found 2 [company] pass-riders and had them sit with/watch over passenger for remainder of flight. Contacted company and asked for flight to be met by the FBI." [58] .

This section has used examples of a number of mitigating factors to illustrate the problems that can arise if investigators are both to assess the potential consequences of an incident and determine what factors combined to preserve the safety of an application. These accounts have been selected because they are each relatively simple. However, the investigators task can become considerably more complicated. For example, the following quotations describe a situation in which the crew of a merchant ship actively intervene to prevent an accident. However, by gathering evidence about the ship and their actions the investigator concludes that their immediate actions had the potential to exacerbate rather than mitigate the incident. This assessment is made even more complex by the fact that the ship and its crew survived both the initial incident and the immediate intervention. The incident began when a load of nickel ore became saturated, settled and started to shift to port.

> "At 2200, or a little before, Padang Hawk suddenly developed a 15 degree list to port. The master, who was in his cabin, immediately went to the bridge and joined the second mate and lookout. The master altered course from 265 degrees to 295 degrees to bring the wind and sea on to the port quarter and reduced the engine revolutions from 110 RPM to 100 RPM... The master decided to ballast starboard side tanks to correct the list. Numbers 3 and 5 starboard topside tanks were filled... At 0145, the master received a reply from the vessels owners advising him to use double bottom tanks to correct the list. The message noted that countering lists by using topside tanks had caused vessels to capsize and it continued: 'Although your vessel is having very high GM due to dense cargo, still high risk of cargo shifting to one side with the roll is high'... The cargo hold bilges were pumped at regular intervals throughout the day. The disposition of ballast was adjusted in accordance with the advice from the owners...
> [Investigators analysis] While recognising the circumstances and the imperative to right the ships list, the master took a significant risk in ballasting the vessel, by adding weight centred high and outboard with an accompanying free surface, without first checking the likely effect on the vessels stability. Although the master was correct in his assessment of the stability, there was a risk of far worse consequences for the vessel and crew, should his intuitive judgement have been faulty. It would have been prudent to use the available resources to calculate the stability of the vessel for all of the conditions prior to transferring any ballast."

This incident begins to illustrate the full complexity involved in both collecting and interpreting evidence about the mitigating factors that influence the development of any incident. The crew intervened in numerous ways to reduce the likelihood that their vessel would be lost. Some of those actions were correct, such as altering the course of the vessel to bring the wind and waves on the port quarter. Other actions were incorrect, most notably the decision to move ballast without first ensuring the stability of the ship. These distinctions reflect what Mackie calls the singular causes of conditions that characterise particular events [508]. Difficulties arise when investigators must move beyond these specific observations to assess the potential severity of an incident without such interventions. Similarly, it is far from simple to determine what might have happen in future situations in which the crew did not perform in the manner described above. One means of reducing this uncertainty, and of supporting other aspects of secondary investigation, is to draw upon evidence from a number of similar incidents.

7.2.3 Identifying Related Incidents

This chapter has described a number of complex tasks that must be performed during the secondary investigation of an adverse occurrence or near miss incident. Many of these tasks are intended to help gather the evidence that will eventually support a causal analysis of 'failure'. Previous sections have argued that ultimately this analysis must look beyond the singular causal factors that contribute to a particular occurrence. Any recommendations should ideally address the more general causes that might lead to similar consequences. In order to do this it is important that investigators gather evidence about similar incidents that may have already occurred. In particular, they must determine whether the singular causes of an adverse occurrence now form part of a wider pattern of failure.

Unfortunately, it can often be difficult to identify common trends in incident reports. Issues of confidentiality and privacy often make organisations reluctant to share information about incidents and accidents. For example, a recent meeting of European air traffic service providers identified a number of common concerns over the impact of that TCAS advisories have upon their ability to sustain safe separations in congested airspace. Aircrews have over-reacted to TCAS warnings; by performing sudden ascents or descents that have infringed on the airspace of other aircraft creating a knock-on effect that can be difficult to counter. Information about a range of similar incidents was passed informally amongst a group of friends from different national providers during a break rather than through any systematic exchange programme. There may of course be information about other similar incidents that is never passed on and so cannot inform the secondary investigation of future adverse occurrences. If anyone is in doubt about this it is instructive to compare the NTSB's report in the collision between a Maryland Rail Commuter and an AMTRAK train [596] with the events leading to the Watford Junction [348] and Ladbroke Grove accidents [351].

A range of further problems prevent investigators from establishing whether an incident forms part of a wider trend. For instance, it can be difficult to ensure that similar events are investigated, analysed and documented in a consistent manner. This is confirmed by both empirical studies and by the more theoretical models of causal analysis. Mackie's notion of a causal field, mentioned above, implies that different investigators may identify different disturbances in the normal state of affairs [508]. This, in turn, can lead them to recognise and diagnose different elements of a causal complex as being salient to a particular incident [508]. Empirical work to back-up this analysis is provided by Lekberg's study of investigator 'biases' [484]. As mentioned in Chapter 3, she showed that different investigators will identify different causal factors within the same incident depending on their previous training and experience. This has profound consequences. If, for example, an investigator were looking for similar incidents in which crew coordination were a causal factor then there is no guarantee that other investigators would have diagnosed this as being significant even if it had indeed taken place. Chapter 10 will introduce a range of analytical techniques that have been proposed to reduce the impact of this problem. For now it is sufficient to understand that such individual differences between investigators may compromise their ability to determine whether or not a particular incident forms part of a more general pattern.

Problems of scale also complicate the task of identifying similar incidents. As mentioned in previous chapters, the ASRS was established in 1976 and now receives an average of more than 2,600 reports per month. The cumulative total is now approaching half a million reports from pilots, air traffic controllers, flight attendants, mechanics etc. Similarly, the FDA's Centre for Devices and Radiological Health's Medical Device Reporting program forms part of a collection of well over 700,000 incidents. Later chapters will introduce a range of innovative technological solutions that are being recruited to support these tasks. In contrast, the remainder of this section looks at a range of more straightforward organisational and managerial techniques that can help investigators to identify similar incidents and common concerns. Fortunately, in some cases it is relatively easy for investigators to determine a pattern of failure. Similar incidents may occur in the same place and within a relatively short-period of time. Under such circumstances, it is readily apparent to many of the individuals who are involved in operating a systems that they may have to address common problems in two or more incidents:

> "Two similar serious incidents were notified to the Air Accidents Investigation Branch (AAIB) at 0630 hrs and 0740 hrs respectively on 6 June 1998, and the investigation

> commenced the same day... The two serious incidents occurred as each aircraft was making an instrument approach to Runway 08 at Ronaldsway Airport, Isle of Man. Both aircraft were using the Isle of Man VHF Omni-Directional Radio Range beacon and associated Distance Measuring Equipment for lateral navigation and distance information respectively. During the course of each of the approaches, each aircraft descended very significantly below the specified descent profile while over the sea to the west of high ground at the Calf of Man and Spanish Head. There was extensive low cloud in the area at these times and in both cases initiation of a climb to avoid possible collision with the high ground occurred once the surface and coastline had been sighted by the pilots involved." [19]

In other cases, organisations may take specific measures to monitor incidents that occur in the same physical location over a more prolonged period of time. This approach has been actively exploited by a number of road traffic management organisations. Sections of road are categorised according to the number and severity of accidents that occur over them in a fixed period of time. Those sections with the worst record are then subjected to an additional level of scrutiny. For example, there may be a detailed analysis of the causal factors behind those incidents that occur on that stretch of road. This analysis and the record of previous incidents help to direct and justify subsequent expenditure on additional safety measures:

> "The junction, near the Lincolnshire Showground, has one of the worst accident records on the A15 between Lincoln and the county boundary. Options for its improvement include a roundabout or staggered junction. It is hoped work could start next financial year. The recommendation for the scheme was made in a safety study commissioned by the Highways Agency in response to the considerable number of road traffic accidents on the A15 in recent years... In the three year period up to 31 May 1998, there were six fatalities on the stretch of A15 covered by the report, 10 serious injuries and 39 slight injuries." [358]

By identifying common causes behind particular incidents, it is possible to justify additional expenditure on more detailed, comparative studies. These investigations might be harded to justify on the basis of individual failures. This approach is exploited by the NTSB. Special investigations are commissioned if investigators identify common causes or consequences in the incidents that they report on. In many instances, these reports simply confirm the initial suspicions that were raised during the initial investigations. However, the additional resources that are invested in these more detailed studies can also reveal more unexpected findings about the potential consequences of a failure. For instance, a recent report demonstrated that cable breakages caused by excavation activities threatened safety in a number of different industries. This was not an unusual finding. However, the potential impact on US air traffic management was not previously appreciated by many other service providers:

> "Network reliability data, compiled since 1993 by NRSC, show that more than half of all facility outages are the result of excavation damage (53 percent), and in more than half of those cases (51 percent), the excavator failed to notify the facility owner or provided inadequate notification... The Federal Aviation Administrations (FAA) study of cable cuts in 1993 documented 1,444 equipment outages or communications service disruptions resulting from 590 cable cuts nationwide over a 2-year period. The majority of cable cuts were related to construction and excavation activities. For 1995, the FAA's National Maintenance Control Center documented cable cuts that affected 32 air traffic control facilities, including five en route control centers. Cable cuts for the first 8 months of 1997 affected air traffic control operations for a total of 158 hours." [598]

Previous quotation have shown how regulators, such as the UK Highways Agency, and investigatory agencies, such as the NTSB, will monitor the common causes and consequences of adverse occurrences. The independent reporting agencies that operate many voluntary reporting systems will also undertake this form of analysis. For example, the ASRS uses three distinct publications to communicate the concerns that are raised within the aviation community. More than 85,000 copies of the CALLBACK newsletter are distributed directly to employees within the aviation community. This includes excerpts from ASRS incident reports with associated editorial comments. It can also contain summaries of ASRS research studies and related aviation safety information. In contrast, DirectLine and the Operation Issues Bulletins are entirely devoted to more sustained investigation about the common causes of adverse occurrences. Although the distinction becomes slightly blurred, the Bulletins cover more immediate concerns whereas DirectLine focuses on incidents that may have arisen over a longer period of time. For instance, the following excerpt shows how DirectLine provides explicit information about common causes, and consequences, in communications failures involving General Aviation (i.e., private pilots):

> "A recent survey of the Aviation Safety Reporting System (ASRS) database on incidents involving General Aviation (GA) aircraft revealed that one third of the GA incidents were associated with communications difficulties... Confusing, erroneous, or misleading statements were the leading type of instructor communications anomaly (30 percent of citations). Delayed or withheld communications by instructors were the next most frequent instructor anomaly (16 percent of citations), and a leading cause of delayed or inappropriate actions on the part of trainees. It is a common technique of flight instructors to allow the trainee to make mistakes in an attempt to develop independent actions and observe the trainee's level of awareness. However, especially during IFR operations, or when compliance with an ATC directive is doubtful, corrective verbal comments by the instructor have a significant impact on flight safety." [228]

Previous sections have argued that investigators gather evidence to help validate their initial hypotheses about the causes of an incident. Information about previous events can provide additional information to guide this validation process. However, there is a danger that beliefs about the causes of a particular incident will be biased by preconceptions about similar incidents. There is also a danger that investigators may diagnose common causes even if two incidents have similar consequences. This is problematic because many different causes can potentially contribute to the same set of outcomes. In spite of these dangers there are, however, considerable benefits if investigators are encouraged to identify common causal factors between similar incidents. This can help to increase the consistency of analysis between investigators. It can help to ensure that similar measures are taken to address the common causes of failure. This, in turn, helps regulatory agencies to determine the success or failure of remedial measures. Such monitoring becomes far more complex if each incident is treated as an individual instance of failure. As a result, many regulatory agencies explicitly encourage these generalisations by publicising common causes and remedies for incidents and accidents:

> "THE PROBLEM:
> Drivers too close to the vehicle in front. 2000 'shunt' type accidents per year on British motorways. Cost of 'shunt' £60 million/year (1989 prices).
> THE SOLUTION:
> Chevron road markings at 40m intervals at problem locations. Signs instructing drivers to keep 2 chevrons from the vehicle in front. Require authorisation.
> THE BENEFITS:
> Study results showed: A reduction of about 15% of drivers 'close-following'. Fewer accidents as driver awareness increased over the site. 56% fewer injury accidents, 89% fewer single vehicle accidents, 40% fewer multiple vehicle accidents, £0.8m/year accident savings (1993 prices). The effect can last at least 18km." [359]

This quotation illustrates how the UK Highways Agency has identified that drivers being too close to the vehicle in front is a common cause in road traffic accidents. They have also gone on to propose chevron road markings as a general solution to this problem and have then gone on to measure the impact of this remedial action. This analysis and the supporting statistics are published in a national compendium of 'techniques and innovative ideas for the better management of the trunk road network' [359]. The success of this document is illustrated by the fact that it has inspired similar initiatives in countries ranging from the Netherlands to Japan. However, there is a danger that such documents will focus the attention of investigators on particular areas of a causal field and that, as a result, on a small subset of possible remedial actions will be taken. This is a particular concern where those remedial actions that are recommended within such a publication are selected for political acceptability rather than effectiveness. Fortunately, the Highways Agency avoids this criticism by publishing statistical evidence to demonstrate the impact of the measures that it advocates. Other organisations have avoided these concerns by adopting a slightly simpler approach. The NTSB does not explicitly identify common causes and general solutions. In contrast, it surveys the recommendations made in incident reports, irrespective of the causes, and then publishes a 'most wanted list'. This, at least publically, avoids any suggestion that all events with particular causal factors can be resolved by the same set of remedial actions. For example, the most wanted safety improvements for highway vehicle occupant protection include the enforcement of state seat-belt laws and an evaluation of whether higher thresholds could safely be allowed for air bag deployment. The corresponding list of commercial truck and bus safety improvements includes general measures to enhance occupant safety, modifications to hours-of-service regulations and higher vehicle maintenance standards.

7.3 Summary

This chapter has focussed on the secondary investigation that, typically, takes place after the primary recipient of an incident report has completed a preliminary report. This phase of an investigation is primarily focussed on securing further evidence about the course of an incident. However, we have argued that this task is guided by a succession of hypotheses about the potential causes of an incident. Evidence is gathered to validate these initial ideas. If necessary, the investigators' causal hypotheses may have to be revised as more evidence becomes available.

It is important to understand some of the distinctions that have been made between the causal factors that contribute to accidents and incidents. For example, Mackie introduced the idea of a causal field, of particular and general causality , of causal complexes [508]. Lewis has pioneered the use of counterfactuals in causal explanations. This work is relevant and significant because it has been integrated into a number of incident analysis techniques that will be introduced in subsequent chapters. Based on this work, we have distinguished between contextual factors, contributory factors and root causes. Contextual Factor are neither necessary not sufficient They are events or conditions that did not directly contribute to the causes of an incident. However, they help to set the scene and establish the context in which an adverse occurrence took place. They may also help to establish that certain factors were NOT significant in the events leading to failure. Contributory Factors are necessary but not sufficient. They are events or conditions that collectively increase the likelihood of an incident but that individually would not lead to an adverse occurrence. These are the 'banal factors' in Reason's observation that "... a detailed examination of the causes of these accidents reveals the insidious concatenation of often relatively banal factors, hardly significant in themselves, but devastating in their combination" [700]. Root causes are both necessary and sufficient. They capture Lewis' notion of causation established by counterfactual reasoning. If a root cause had not occurred in the singular causes of an incident then the incident would not have occurred.

Later sections went on to examine sources of evidence that can be used to identify contextual factors, contributory factors and root causes. It was argued that the use of independent expert witnesses can help to combat the natural biases that can persuade investigators to favour particular causal hypotheses. However, there is also a danger that such witnesses may themselves be biased. In order to address this problem, we developed Spohrer and Maciejewski's [754] ten commandments for Chemists acting as expert witnesses during criminal investigations. We presented a more general set of guidelines based on these heuristics so that they might support the wide range of experts who are called upon to support incident investigations.

Evidence about the causes of an incident can also be extracted from the automatic monitoring devices whose logs are preserved during the initial response to an incident. However, this chapter has reviewed the considerable managerial and technical problems that continue to affect the use of these critical data sources in many industries. For example, it can be difficult to ensure that these devices are correctly maintained. There have also been instances where monitoring devices are incorrectly configured to the individual standards that many commercial organisations are creating. Even if data is correctly recorded, problems can arise when duplicating data or in finding a correctly configured reader. Although many of these problems are being addressed both by regulators and manufacturers, they continue to be document in incident reports that lament the lack of automated logs.

The second half of this chapter focussed on the importance of gathering evidence about the consequences, as well as the causes of adverse occurrences. In some situations this can be relatively straightforward. The effects of any failure can be directly witnessed by those involved in the immediate precursors to an incident. In other contexts, the individuals who contribute to a failure may have no idea of the impact that their actions have had. For example in transportation systems, problems can occur many miles away from the maintenance facility that contributed to the failure. In medical systems, the consequences of an incident may not manifest themselves until years later when the patient's physical well-being and quality of life may be seriously compromised.

The problems of gathering evidence about the consequences of an incident are further complicated by the fact that investigators may have to account for mitigating factors. These interventions can reduce the consequences of a particular incident. As a result, investigators may choose to treat the occurrence as if the intervention had not taken place. This approach exploits the notion of a worst plausible outcome . However, a limitation with this technique is that it can be difficult to predict the ways in which apparently trivial failures can quickly escalate into major accidents. Further problems are created by the difficulty of establishing possible combinations of contributory and mitigating factors that are likely during any future failure.

One means of addressing the uncertainty that arises during the secondary analysis of any incident is to gather as much information as possible about similar incidents. This can be done by investigating records of previous failure in the same location or within a similar period of time. It can also be done by examining regulatory and investigatory 'hit lists' of common causal factors in adverse occurrences. Incident reporting systems may also provid information about previous problems. These alternative sources of evidence help to increase the investigators confidence in any generalisations that may be made about the causes and consequences of particular incidents. However, there is also a danger that they may inadvertently bias any investigation towards the findings of previous investigations. Rather than looking at each incident as a potentially unique occurrence, investigators might simply attempt to place it within pre-existing categories of superficially similar incidents.

To summarise, this chapter has stressed the importance of gathering evidence about the causes and consequences of adverse occurrences. It has also explained why it can be so difficult to achieve this. Technical difficulties continue to frustrate automated logging techniques. The problems of determining a plausible worst case scenario frustrate attempts to gather evidence about possible consequences of previous incidents. The following chapters, therefore, present a range of techniques that are intended to address these problems.

Chapter 8

Computer-Based Simulation

The previous chapter identified the main activities that must be conducted during the secondary investigation of any incident report. These, typically, focus on gathering evidence to both inform and validate initial hypotheses about the causes of an adverse occurrence. This chapter, in contrast, looks at one aspect of this validation process. It seems clear that any causal hypothesis must be consistent with what we know about the course of any incident. Support for such assertions is often provided by simulation and reconstruction techniques. The Rand report into the National Transportation Safety Board's (NTSB) investigation techniques emphasised this in their overview:

> "When a complex system fails, the number of potential scenarios rises proportionately. NTSB investigators must carefully unravel the performance of many highly integrated systems, a very time-consuming task requiring a diverse set of skills. Often, this requires extensive and costly salvage and reconstruction of the aircraft. Complexity affects more than just staff workload. The growing complexity of aircraft crashes also has a profound effect on how investigations must be structured to reveal hidden failure modes." [482]

This quotation reveals the dual nature of reconstruction in many modern incident investigations. Firstly, reconstruction involves the rebuilding of components and sub-components to identify causal information from the physical damage that often occurs during major incidents. Secondly, there is the more abstract notion of event simulation in which investigators piece together the more complex causes of an incident drawing upon the physical evidence and also from the other forms of evidence that are gathered during a secondary investigation. This might suggest a firm distinction between physical reconstructions and virtual simulations. The terms 'simulation' and 'reconstruction' are, however, often used inter-changeably. As we shall see, this ambiguity can partly be justified by the way in which limited physical reconstructions are being used to provide the data that drives more general computer-based simulations.

8.1 Why Bother with Reconstruction?

The term 'reconstruction' has traditionally been used to describe the way in which physical evidence is re-assembled to provide clues about the sequence of events leading to failure. For example, the US Army's accident and incident investigation guidelines constraint the following recommendations for the analysis of rotor or propeller failures. As can be seen there is a requirement to reconstruct the entire assembly if at all possible:

1. Collect and inventory; reconstruct the whole assembly if possible.
2. Examine damage / scarring to determine if systems were turning at impact and if power was applied at impact.
3. Examine all linkage from cockpit controls to systems for continuity/disconnect, all bearing assemblies and / or blade grips for failure prior to impact.

4. Check for serial numbers of blades / propellers against historical records. [806]

A number of published guidelines provide detailed information about the ways in which such physical reconstruction should be conducted [750]. Much of this information varies from domain to domain. For instance, the construction and operational stresses of aircraft components are quite different from those relating to automotive components. Klepacki, Morin and Schaeffer's guidelines for evaluating post-incident flight control trim system configurations are highly domain specific [448]. There are, however, some similarities between the techniques that are used to support incident reconstruction in several different industries. For instance, the techniques for establishing the velocity and angle of impact damage show strong similarities across several different domains [869].

The opening paragraphs identified two forms of reconstruction. The first focussed on the physical rebuilding of damaged components to gain further information about the failures that contributed to an incident or forces that arose in its aftermath. The second aspect of reconstruction deals with the way in which information is used to describe the course of events over time. This centres on the process of assembling fragmentary evidence to produce a coherent account of an adverse occurrence. For example, the US Air Force requires that investigators reconstruct the sequence of events that leads to an incident [794]. They must map route segments. They should provide a vertical view of maneuvers. The account may include artists conceptions or models to explain the course of events. The intention is to "explain what the plan was, what should have happened if things had gone right, who was in charge, what were the rules of engagement and were they followed, where things went wrong, what should the aircrew have done, and what were the aircraft parameters at ejection or aircraft impact" [794]. As can be seen, generic requirements are specified together with more domain specific guidelines that relate narrowly to aviation accidents.

The wealth of guidance on the physical reconstruction of safety-critical systems is not matched by similar sources of advice on the reconstruction of events leading to incidents and accidents. As a result, the remainder of this chapter focuses on techniques that can be used to build coherent models that explain how different events contribute to an adverse occurrence. It is important not to underestimate the importance of these reconstructions. They are intended to produce a coherent account of the course of an incident from many disparate pieces of evidence. In other words, they are intended to explain *what* happened while causal techniques present *why* it happened.

There are many different ways in which to build these event reconstructions. In some domains, it is also possible to stage physical reconstructions. These re-enactments are often used by the Police to trigger witness recollections and elicit further information about an incident. However, there are obvious limitations with this approach. For example, such reconstructions can expose individuals to further danger. There are ethical considerations involved in re-enacting a failure in a working foundry or chemical plant. There is also the danger, described in Chapter 5 that such 'realistic' reconstructions may trigger further psychological problems for those involved in an incident. It may even trigger false memory syndrome in some cases. Fortunately, a range of alternative techniques can also support the reconstruction of safety-critical incidents. For example, computer-based simulations enable investigators to step-through the events leading to an incident using three dimensional animations. The second half of this chapter identifies a number of these interactive tools. There are, however, a number of limitations with computer-based simulations. For example, highly interactive models provide a good impression of the events leading to component failure. Unfortunately, they cannot easily be used to recreate the events leading to managerial or regulatory failure. I have yet to see virtual reality simulations recreate a board meeting or a management conference in which safety investments were turned down!

Fortunately, a range of graphical and textual notations can be used to avoid such limitations. They can be used to sketch the events leading to an incident at a far greater level of abstraction and can, therefore, also capture events leading to managerial and regulatory failure. Many of these notations provide inference techniques that can be used by investigators to clearly distinguish what can, and what cannot, be concluded from the evidence that is assembled. These reasoning techniques can also be used to identify inconsistencies and omissions in incident reconstructions. The following chapter focuses on these formal and semi formal notations for event reconstruction. The second half of this chapter reviews computer-based modelling techniques. It is important to emphasise, however, that these modelling notations and the computer-based simulations mentioned above are *not* primarily intended as tools to support the causal analysis of adverse occurrences. They are intended to reconstruct the course of events that contribute to an incident. In contrast, Chapter 10 presents techniques that help to distinguish root causes and contributory factors from the contextual information that can be represented in a formal model or an interactive simulation.

Some formal and semi-formal modelling techniques are supported by computer-based tools. It is possible to derive interactive computer-based simulations from abstract notations. As a result, the previous distinctions between computer-based simulation and abstract modelling can become blurred. However, these distinctions are retained because they are useful in distinguishing between different sets of concerns that affect what can be complementary approaches. For instance, formal mathematically-based notations can be difficult to interpret by non-mathematicians. In contrast, interactive simulations can often lead to unwarranted interpretations if analysts are seduced by particular animation techniques. For now it is sufficient to emphasise that natural language provides the most accessible and widespread medium for building reconstructions. The following quotation illustrates how most organisations summarise the events leading to an incident These accounts are used to provide investigators and managers with a common ppoint of reference during any investigation. They are gradually refined as additional evidence is obtained until they are eventually integrated into a final report:

> "About 10 p.m., unknown to the controller, the pipeline ruptured at a location near Gramercy, Louisiana. At 10:01:53 p.m., the supervisory control and data acquisition (SCADA2) system reported high-pump-case pressure at Garyville. The SCADA system activated an audible alarm and also displayed a message on a display screen. Almost immediately, the SCADA system sounded and displayed alarms reporting that certain pumping units at the Garyville station had automatically shut down because of low suction pressure (low liquid pressure on the inlet side of the pump). At 10:02:30 p.m., the SCADA system reported a line balance alarm." [593]

This quotation shows the way in which natural language can represent the events that occurred immediately before the rupture. The model of this incident is constructed around a number of key incidents, such as the first SCADA2 system report, for which the timing is known. It is important to note that these proximal events cannot be viewed as catalytic because no argument is provided to demonstrate that they directly *caused* the incident. Of course, it is also possible to use natural language to describe the more distal, latent causes of this failure. This again illustrates how the reconstruction of an incident is guided by implicit causal hypotheses. In this case, evidence about previous excavations near the Marathon pipeline suggests that damage might have been caused during these earlier operations:

> "The investigation determined that in 1995, LaRoche Industries, Inc., arranged for excavation of and repairs to various portions of its 8-inch pipeline, which was located about 30 feet from the Marathon pipeline. These excavations took place in September and October 1995 in the vicinity of the Marathon pipeline rupture... According to officials from LaRoche's contractor, the equipment operators were told by LaRoche superintendents that no pipelines were located in the area of the Marathon pipeline." [593]

The previous quotations illustrate how prose descriptions can be used to draw upon various sources of evidence in order to reconstruct the events that led to an adverse occurrence. These natural language accounts must consider the many different factors that contribute to the increasingly complex incidents that have been described in previous chapters. The following paragraphs briefly summarise the types of information that must be captured. These can be loosely categorised as belonging to three distinct stages in an 'incident sequence'.

Initial Conditions

The initial conditions describe the normal operating state of the system and its environment before an incident occurs. The following quotation illustrates how NTSB investigators describe the initial state of the system as part of a passage that sets the scene for the failures that follow:

> "On May 23, 1996, a pipeline controller was on duty in Marathon Pipe Line Company's pipeline operations center in Findlay, Ohio, operating and monitoring a 68-mile-long segment of Marathon pipeline located in Louisiana. This pipeline is used to transport hazardous liquids between a refinery at Garyville, Louisiana, and a station at Zachary, Louisiana. Pumps at the Garyville refinery pressurise the pipeline and generate the power to transport the liquids to the Zachary station. About 9:53 p.m. central daylight time on May 23, the pipeline controller had just completed operations to transport a batch of unleaded gasoline through the pipeline. He then remotely executed commands to introduce into the pipeline (behind the gasoline) a batch of 125,000 barrels of low-sulfur diesel fuel." [593]

This excerpt establishes the general topology of the pipe network. It also introduces the controllers' tasks immediately before the incident took place. As can be seen, the initial conditions in a reconstruction describe the situation as it existed before any adverse incidents occurred. This introductory section closes at the moment when the remote command was executed. At this point the pipeline ruptures and the reconstruction continues with the first of the quotations cited in the previous section.

It is important to emphasise that the initial conditions that are described by a reconstruction need not be normative. They may not satisfy relevant safety regulations or recommended operating practices. In particular, the initial description of the system might indicate that there were frequent safety violations during normal operation. For example, the NTSB overview of another pipe rupture describes how the company "had procedures in place at the time of the accident that were applicable to general construction activities in proximity to its pipelines, but it did not have procedures specific to directional drilling operations' [597]. These operating practices did not have adverse consequences until the company attempted to install a distribution main parallel to a gas transmission pipe. The proximity of the installation damaged the transmission pipe and this led to a rupture that the NTSB estimates cost in excess of $2 million.

Catalytic Failures

Reconstructions must also describe the events that helped to move the system from its initial condition towards an eventual incident. These events include the catalytic failures that are often central to any subsequent investigation. In many instances these are clear cut. For example, some pipeline ruptures can be directly related to specific (catastrophic) events:

> "About 4:50 a.m. on October 23, 1996, in Tiger Pass, Louisiana, the crew of a dredge dropped a stern spud into the bottom of the channel in preparation for dredging operations. The spud struck and ruptured a 12-inch-diameter submerged natural gas steel pipeline owned by Tennessee Gas Pipeline Company. The pressurised natural gas released from the pipeline enveloped the stern of the dredge and an accompanying tug, then ignited, destroying the dredge and the tug." [592]

It is important again to reinforce the point that such events do not provide any direct explanation of the root causes of an incident. Such underlying causes are often embedded within the initial conditions that were mentioned in the previous paragraph. These conditions combine to create a situation in which catalytic events have the potential to trigger an incident or accident. For example, the dredging operation had to rely upon the gas company's practices and procedures for locating, marking, and maintaining markers for gas pipelines through navigable waterways. The potential for a catalytic event was also created by the lack of Federal requirements for placing and maintaining permanent markers where gas pipelines cross navigable waterways.

It should also be emphasised that there are other incidents in which it is far harder to identify the catalytic events that actually trigger a failure. For instance, in the incident in which directional drilling was cited as the root cause of the pipeline fracture, it is unclear as to which aspect of the operation actually *caused* the failure. The rupture occurred as the distribution pipeline was being returned to full service and not as an immediate consequence of a catastrophic operation as was the case in the dredging incident. Some evidence pointed to the fact that a reaming tool left gouge marks in the vicinity of the rupture but it is difficult to be certain of the precise operation that resulted in the eventual failure. In such circumstances, there must be some ambiguity or uncertainty in an eventual reconstruction of the catalytic events that contributed to an incident.

Liveness Conditions

There are, typically, several moments when operators or automated systems might have intervened in order to prevent an incident. 'Liveness' conditions, therefore, not only describe the catalytic events introduced in the previous paragraph. They also reconstruct the manner in which these safeguards failed. In other words, liveness conditions also describe the events that enabled an incident to progress towards its final consequences. For instance, the NTSB account of the gasoline release describes how operators failed to detect and respond to the initial alarms from the SCADA system. Had they responded to these warnings sooner then the full consequences of the incident might have been considerably reduced:

> "The pipeline controller continued to receive alarms. Initially, he acknowledged each one individually, but believing that each subsequent alarm was related to operations at the refinery, he elected to simultaneously acknowledge all the alarms and the alarm text messages without attending to the nature of each alarm... The controller said he called Garyville and discussed the situation with the station operator there. The station operator confirmed the automatic pump shutdowns. The station operator determined that the Garyville refinery was, indeed, loading product to a barge. Even though refinery personnel reported that the volume of product being delivered was insufficient to have caused the SCADA system to alarm, the pipeline controller and the station operator concluded that the loading of the barge had precipitated the alarms and the pump shutdowns." [593]

To summarise, reconstructions must represent the initial conditions or context for an incident. They must also describe the way in which catalytic failures initiate the events leading to failure and how liveness conditions create the necessary conditions for an incident to develop. The following section describes the final component of any reconstruction; the events that take place in the aftermath of an adverse occurrence.

Consequences

The reconstruction of an incident cannot simply stop at the point with catalytic events. It is a truism that more lives are lost through failures in the 'golden hour' after an accident has occurred than are killed by the catalytic events themselves. Chapter 7 has emphasised the importance of incident reporting as a means of assessing an organisations ability to respond to or mitigate the adverse consequences of any failure. In consequences, reconstructions must go beyond the catalytic events that are often the focus of attention in the aftermath of any incident. This point can be reinforced by the consequences of two further pipeline failures. On October 30th 1998, excavation work damaged a 24-inch diameter gas main in Chicago. This released natural gas that ignited about forty minutes after the initial rupture. The immediate consequences of the failure included substantial damage to a high-rise block of appartments. However, the prompt response to fire and police personnel completely evacuated the building so that no-one was injured. In contrast, a similar incident two months later left four dead, one person seriously injured and ten people, including two firefighters and a police officer, with minor injuries:

> "An engine company with a lieutenant and three firefighters arrived within minutes of fire department notification. Firefighters attempted to take gas concentration readings with a gas monitor, but the monitor had not been calibrated in fresh air and gave invalid or unreliable readings. Firefighters continued to attempt readings with the improperly calibrated instrument, all the while working in an environment in which they described the gas smell as pretty bad. At no point did firefighters check buildings near the leak site to determine if natural gas was accumulating or to help assess the need for

> a possible evacuation, even though the gas line was continuing to release gas that could migrate through the ground and into nearby buildings, where it could present a danger of explosion. Two of the firefighters near the leak site returned to their truck as soon as two gas company employees arrived. It should have been obvious to the firefighters that a threat continued to exist and that the situation could worsen. The Safety Board therefore concludes that firefighters of the St. Cloud Fire Department responded quickly to the scene of the leak; however, once on the scene, the firefighters actions did not fully address the risk to people and property posed by the leak or reduce the consequences of a possible fire or explosion." [602]

Many investigation authorities have placed increasing emphasis on response time targets for emergency services. This incident again illustrates that a prompt response must be backed up by effective actions if we are to mitigate the effects of such incidents. It is, therefore, necessary for reconstructions to explicitly represent the actions that are taken in the aftermath of a catalytic failure.

Previous paragraphs have explained how it is important to reconstruct the initial conditions that create the context for any incident. It has also been argued that investigators must produce a coherent account of the catalytic failures that trigger those events that lead to an adverse occurrence. Liveness conditions must also be reconstructed. These represent the way in which defences must be breached and warnings ignored in order for a catalytic event to escalate into a major failure. Finally, it has been argued that reconstructions must also consider the consequences of any incident. These are partly shaped by the nature of the failure but also by the interventions that help to mitigate those consequences. However, there are relatively few benefits to be obtained from simply developing accounts that describe how all of these events occurred during the course of an incident. The following paragraphs, therefore, identify ways in which reconstructions can be used to inform the investigation of an adverse occurrence and, ultimately, to reduce the likelihood of any recurrence.

8.1.1 Coordination

Chapter 7 has described how incident investigations draw upon the the work of many different experts. Forensic scientists, metallurgists, meteorologists, software and systems engineers as well as human factors experts all contribute to these enquiries. It can be difficult to coordinate the activities of these different group. There is a risk that necessary tasks may be omitted or needlessly duplicated. It is, therefore, important that investigators have some means of monitoring and coordinating the finite resources that they can deploy to support their enquiries into an incident. Reconstructions provide a useful tool to support these managerial tasks. They provide a model of the events leading to an incident. Individuals with different domain expertise contribute to different aspects of these reconstruction. For example, metallurgists can describe the conditions that might have contributed to catalytic metal fatigue. Human factors experts can identify salient events in a crews' response to an incident. These different contributions must be pieced together to form a coherent view of a complex incident.

There are a number of ways in which experts contribute to the overall process of incident investigation through their participation in any reconstruction. These can be summarised as follows:

- *broadening the causal field.* One of the key roles for any expert is to help broaden the causal field of any analysis. Chapter 7 explained how this field represents a subjective frame of reference that individuals or organisations use when trying to explain what has happened in a particular situation. If an event does not have an impact upon the causal field then it may not be identified as playing a significant role in the course of an incident. Prior expertise plays a significant role in knowing where to find the evidence that indicates certain events have taken place. Without this expertise, evidence might not be found and a reconstruction might not include necessary information about an incident. .

- *determining the salience of events.* In contrast to the experts' role in broadening a causal field, they may also identify certain events as not playing a significant role in a particular incident. These events might then be omitted from any subsequent reconstruction. There is, of course, a potential danger in this if those events later emerge as having a more important role in course of events. It is, therefore, important to document the reasons for such omissions. It is also important to stress that reconstructions support but do not replace causal analysis. For example, it is often impossible for any single expert to diagnose the root cause of an incident without referring to the work of their colleagues. In consequence, investigators use reconstructions to provide an overview of the evidence that is collected about the diverse events that lead to an incident. This overview of events must then be interpreted and analysed to distinguish between contextual factors, contributory factors and root causes. Techniques that support this causal analysis will be described in the next chapter. In contrast, this chapter focuses on techniques that can be used to reconstruct the 'flow' of events leading to an incident and the consequences that stem from such failures.

- *determining knock-on effects of events.* By participating in any reconstruction, experts are also forced to consider the ways in which events in other areas of a system can affect their area of expertise. For instance, a human factors expert must consider the impact of prevailing weather conditions if a meteorological experts have indicated that this may be a factor. Conversely, building a reconstruction can also help investigators to identify the knock-on consequences that particular events will have throughout a system. This is a by-product to the tasks involved in developing a narrative account that links together the evidence that is available in the aftermath of an incident.

- *eliminating particular events.* The development of a reconstruction can force analysts to determine whether or not particular events actually did contribute to an incident. For example, the Minnesota pipeline investigation initially questioned whether the location of the line had been incorrectly marked out. Evidence had to be provided to determine whether this was a potential problem before any detailed reconstruction could be built. If it had been incorrectly marked out then additional resources would have been deployed to examining the events that contributed to this failure. However, the NTSB investigation determined that "the marked location of the ruptured gas line was accurate and therefore not a factor in this incident" [602]. As a result, the prose description of the incident does not focus in great detail on the initial surveying of the line.

- *forcing the resolution of inconsistency.* As investigators contribute to the development of a reconstruction, it is likely that a number of inconsistencies and omissions will be identified in the overall timing of events. As we shall see, other anomalies can arise. For example, eye-witness testimonies often place the same individual at two different locations at the same moment in time. Such inconsistencies can be resolved by finding evidence to discount one statement. Alternatively, two or more alternative reconstructions can be developed to explore several different incident scenarios. However, contradictory witness statements are not the only source of inconsistency in incident reconstructions. Other problems relate to the group processes that affect incident investigation teams. As mentioned, these enquiries often involve heterogeneous teams of domain experts. The members of these groups often have different backgrounds and training. Partly as a result of this, the conclusions of one analyst about the probable ordering of events need not accord with those of another. For example, the Fire Service and Ambulance accounts of the Clapham rail crash differ in several important respects [502]. If these problems are not resolved during the reconstruction phase then there is a danger that any causal analysis will be jeopardised because of contradictions in the evidence that it relies on. There is also the danger that the eventual incident report will contain inconsistent information about the sequencing of events leading to and stemming from catalytic failures.

A number of important consequences stem from this use of reconstructions as a means of coordinating the various activities that contribute to an incident investigation. In particular, natural language descriptions, interactive computer simulations or other diagrammatic techniques must be capable of capturing the key events identified by different domain experts. It is also important that those domain experts can read and understand the resulting reconstructions if they are to validate the models that are produced. This is a non-trivial requirement. The complexity of many incidents makes it difficult to trace the ways in which system 'failure' and operator 'error' interact over time. For example, many incident reports now run to several hundred pages of prose narrative.

8.1.2 Generalisation

Incident investigations are intended to determine what caused a failure and to identify means of preventing any recurrence. As we have seen, reconstructions play an important role in validating the evidence that, in turn, supports subsequent causal analyses. They also play an important role in identifying ways in which an incident can recur. The process of identifying those events that contributed to a particular incident helps to inform subsequent investigations to determine whether those events might recur in isolation or in combination with other failures. In other words, in order to identify the general causes of future incidents it is important to understand the causes of the particular incident under investigation [677]. Those causes can only be accurately established by ensuring the validity of any reconstruction.

An important application of reconstruction techniques is in the development of training scenarios. These, typically, start with the events that lead to previous failures. For example, the following citation comes from an NTSB incident report that explicitly considered the ways in which reconstructions of previous incidents were used by some utility companies to drive simulation-based training:

> "The UGI's emergency plan requires each employee who is responsible for responding to emergencies to participate in annual simulation board exercises. Each exercise is prepared by the UGI's distribution engineering personnel and includes scenarios about a system shutdown or loss of a major gas supply line, a shutdown or loss of a district regulator station, or a major line break within the distribution network. The scenario may be based on previous incidents or on incidents described in Safety Board reports. Each exercise must include a step-by-step analysis of the procedures for investigating, pinpointing, and repairing leaks and of the procedures for taking emergency actions and protecting people and property." [588]

These simulations can be used to determine how well teams can cope with the situations that previously confronted their colleagues. However, crews seldom intervene in exactly the same manner as their colleagues. As a result, their actions help to shape new scenarios that differ from the events that occurred in the original incident. Simulation based training, therefore, enables crews to explore more general forms of failure that are based on the particulars of a singular incident. This close relationship between training and reconstruction is emphasised in the Rand report into the NTSB:

> "The NTSB should review its internal technical capabilities to support future accident investigations, including the potential for crash reconstruction and the requirements for system testing in support of complex accident investigations. The safety boards long-term requirements for facilities should include consideration of their use for staff training, recognizing that facilities can serve a dual function." [482]

It is to be hoped that future incidents will not occur in exactly the same way as previous incidents. If they do then this clearly indicates the failure of a reporting system to address the underlying causes of any failure. However, it is not clear how the particular details of an adverse occurrence can be used to anticipate other, more general forms of future failure. The following list identifies a number of ways in which reconstructions can be manipulated to support this form of analysis. The intention is to manipulate the reconstruction in order to either identify training scenarios or to ensure that any recommendations address a wide range of potential future failures:

- *transposition of events.* The most obvious way of generating alternative incident scenarios from any narrative of a particular incident is to alter the sequence of particular failures. For example, in the Garyville incident mentioned in previous sections, a reconstruction might simulate the rupture before, during or after the completion of the controller's transportation command on the batch of unleaded diesel. It is important to stress that the undirected transposition of events will not always lead to failure scenarios. It can also lead to scenarios that might seem extremely implausible. For instance, it seems unlikely that the supervisory control and data acquisition (SCADA) system might generate the high-pump pressure alarm before the pipe failure event. The irony is that many engineers and designers have failed to adequately account for those scenarios that were dismissed as implausible before they occurred [65].

- *omission of adverse events.* A further means of generating alternative scenarios is to omit some of the failures that arose during a previous failure. This can simplify the demands that a training scenario may place upon system operators. Additional complexity can be gradually introduced as teams become more skilled in responding to an adverse situation. This exploits the 'training wheel' approach in which supports are gradually removed from operators as their confidence grows [155]. A particular benefit of this approach is that it can be used to prioritize the allocation of resources to improve system defences. For example, the NTSB report into Minnesota explosion hypothesised that "had the gas line in this accident been equipped with an excess flow valve, the valve may have closed after the pipeline ruptured and the explosion may not have occurred" [602]. This assertion can be tested both using laboratory simulations of the gas flow within the system. Operator performance can also be assessed by reconstructing the course of an incident as though this defence had existed. If the crew can consistently respond to correct and mitigate these alternative scenarios then the proposed defences can be shown to offer some protection. If crews cannot mitigate a failure with these defences then they are unlikely to provide sufficient protection.

- *exacerbation of adverse events.* Reconstructions not only provide scenarios that can be used to assess the effectiveness of potential defences, they can also be used to assess the consequences that may ensure if existing defences are compromised. As we have seen, one of the key differences between incidents and accidents is that particular safety features intervene to mitigate the consequences of failure. We have already argued that an important component of any incident investigation is to determine the 'worst plausible outcome' . These scenarios are again critical both in guiding training and in assessing the potential effectiveness of any remedial actions. For example, the following citation illustrates how NTSB investigators often consider the circumstances that might have exacerbated any failure:

  > "In this accident, the speed and extent of the gas release and fire placed all crew-members aboard the dredging vessels in grave danger. Fortunately, despite the early hour, most crewmembers were awake, alert, and able to respond quickly to the emergency. Given the rapid ignition of the natural gas and the extent of the damage to the vessels, had this accident occurred while most of the crew was sleeping, numerous serious injuries or fatalities may have occurred. The Safety Board concludes that in even a slightly more serious accident, Beans emergency procedures, because they did not require that a precise count be kept of the number of personnel on board the companys vessels at all times, would have been inadequate to account for and facilitate the rescue of missing crewmembers, increasing their risk of serious injury or death." [592]

  Compound simulation techniques provide another means of preparing for plausible worst case scenarios. This approach combines elements of one incident with events that occurred during another previous failure. The result is to create hybrid incidents that blend multiple problems identified during previous incidents. This approach is motivated by Reason's plea not to consider failures in isolation [701].

There are a number of further problems that affect the generalised use of simulations to investigate potential failures. In particular, it is difficult to accurately reproduce operator behaviours under 'experimental' conditions. However, such caveats have to be balanced against the benefits that reconstructions provide in generating the 'what if' hypotheses that direct future development.

8.1.3 Resolving Ambiguity

A key benefit of reconstruction is that it helps investigators to identify omissions and inconsistencies in the evidence that they gather about an incident or accident. This can be illustrated by the NTSB report into the St. Cloud pipeline failure:

> "At about 10:50 a.m. on December 11, 1998, while attempting to install a utility pole support anchor in a city sidewalk in St. Cloud, Minnesota, a communications network installation crew struck and ruptured an underground, 1-inch-diameter, high-pressure plastic gas service pipeline, thereby precipitating a natural gas leak. About 39 minutes later, while utility workers and emergency response personnel were taking preliminary precautions and assessing the situation, an explosion occurred. " [602]

This high-level summary is typical of the sparse information that may be available in the immediate aftermath of an incident. Lack of evidence can prevent investigators from building more detailed reconstructions of the events that contributed to a failure. However, it is possible to use such prose descriptions to help target those events that deserve closer scrutiny. One technique is to scrutinize these narratives in order to identify any ambiguities that require further clarification. These ambiguities partly stem from the flexible ways in which investigators can use natural language to support a number of different interpretations based on the same sentence. For example, the previous quotation includes the observations that the pipe was ruptured by the crew 'while attempting to install a utility pole support anchor in a city sidewalk...'. This abstract description could refer to any number of more detailed procedures that the crew could have been performing in order to achieve their goal of installing the utility pole. They might have been drilling, using a sledgehammer to break the sidewalk, using an auger to secure the anchor etc. If the exact operation that was being performed at the moment of the rupture was critical for a more detailed understanding of the course of events, as is likely to be the case, then investigators must gather more detailed evidence about the crews' actions. The following list, therefore, identifies a number of different forms of ambiguity that can occur in natural language reconstructions, or accounts, of safety-critical incidents:

- *ambiguity of time.* The previous account referred to real-time, '10:50am' and 'About 39 minutes later...'. It also used less precise relative timings that are implicit in phrases such as 'while attempting', 'thereby precipitated'. An important strength of such descriptions during the initial stages of investigation is that it is possible to construct models that describe several different real-time orderings for the events that are identified. For example, the phrase 'About 39 minutes later...' describes possible reconstructions in which the explosion occurred at 38, 39 or 40 minutes after the initial rupture. The scope of the interval is only bounded by the readers' interpretation of 'about 39 minutes'. These slightly vague timings can be made more concrete as further evidence is obtained. However, there are also examples where exact timings cannot ever be confirmed. For example, the time-line of events are often incomplete [588]: Alternatively, if the timing information is not considered significant to the overall analysis investigators may deliberately choose not to expend finite resources in resolving such ambiguity. All of this contrasts sharply with many of the computer based simulations that we shall explore

Time	Event
6:48 p.m.	The EPAI foreman called the home of the EPAI Vice President.
6:?? p.m.	The foreman instructed his crew to trace the gas line back toward Utica Street to shut off the gas valve.
6:50 p.m.	The EPAI foreman called the UGI emergency telephone number, advising that they definitely hit the gas line and broke it.

Table 8.1: Excerpt from the Incomplete Time-line of a Gas Explosion

in subsequent sections, these typically require that investigators commit themselves to precise intervals in which events can occur.

- *ambiguity of place.* The previous account only provides a high level view of the events that contributed to the incident. As we have seen, the US Air Force requires that investigators provide maps of the relative movements of aircraft during an incident. The use of terms such as 'in a city sidewalk in St. Cloud' provide an insufficient level of detail for most investigations. Clearly, any secondary investigation would be expected to produce a more detailed survey of the incident. This illustrates the important observation that any reconstruction will, typically, have to exploit a variety of media if it is provide a complete overview of the many different sorts of information that must support any subsequent causal analysis. Increasingly this may include video footage as well as graphical sketched and textual accounts.

- *ambiguity of action.* The previous summary uses natural language to provide a high-level view of the events leading to the incident. As mentioned, this use of prose provides considerable benefits in terms of flexibility and comprehension. It supports multiple interpretations when necessary evidence is not available. Additional details can be introduced as they are gathered. These comments not only apply to the representation of time and place, it also refers to the account of the crews' actions. Phrases such as 'while utility workers and emergency response personnel were taking preliminary precautions' provide few insights into their actions. Again, evidence must be gathered to determine whether or not their precautions had a significant impact upon subsequent events. There are further benefits of ambiguity in the representation of actions. For example, it is possible to indicate that a crew member performed certain tasks without describing the components, or sub-tasks, that this might have involved. This provides significant benefits if, for instance, these components can be understood from the context of the actions. Problems will, of course, arise when other members of the investigation team do not have the necessary domain knowledge to interpret what this task might have involved. It may also cause problems if, in fact, necessary sub-tasks were either omitted, duplicated or interrupted. Such complexities are masked by this ambiguous action description.

- *ambiguity of motivation.* The previous account provides little information about the potential factors that motivated the crew's decision to anchor the utility pole in that particular location. As with the other forms of ambiguity; there are multiple reasons why natural language descriptions avoid spelling out such factors. In the aftermath of an incident, it can be very difficult to gather objective evidence to support explicit interpretations of individual performance. It is also the case that many investigators lack the human factors training to be confident in proposing more explicit models of the cognitive and perceptual factors that influence operator behaviour. Most reconstructions entirely avoid representing or reasoning about the internal cognitive factors that motivate particular actions. Both natural language descriptions and computer-based simulation techniques, typically, therefore, focus on observable actions only.

- *ambiguity of cause.* The previous description is ambiguous about what exactly caused the incident. It might have been caused by the gas service provider failing to document the position of its pipeline. It might also have been caused by mistakes in siting the anchor for the utility pole. Although these both contributed to this singular incident, it is unclear whether either is necessary and sufficient in the general case. Partly as a result of this causal ambiguity, many investigation agencies deliberately separate the process of finding out *what* happened from explaining *why* it occurred [423]. We have, however, argued that these activities are strongly linked. Reconstructions help to validate and guide causal hypotheses. Later sections will argue that it is, therefore, extremely important that tools and techniques be provided to link these two complementary activities. In particular, it is important that causal ambiguities should not be left in a final report so that the reader is left in considerable doubt about the root causes of an incident.

The previous paragraphs have tried to emphasise that there are often good reasons for ambiguity in the initial reconstruction of an incident. For example, temporal ambiguity can occur because there may not be sufficient evidence to determine the exact moment at which an event occurred. Even in the later stages of reconstruction, ambiguity still plays an important role in the communication of information about complex failures. For example, ambiguity of action can help to abstract away from the exact sub-tasks that an operator or system performed if those sub-tasks can be assumed from the surrounding context of the description and those sub-tasks did not play a significant role in the course of an incident. The following paragraphs summarise several of these reasons why ambiguities may remain in reconstructions of the events leading to failure.

As mentioned, there may not be the evidence available to provide definitive information about the specific course of events leading to an incident. The following synopsis illustrates how in some situations it is only possible to gather superficial facts about the course of an incident. This incident involved a relatively small business jet. Without an advanced flight data recorder or detailed information about the pilot's actions it is difficult to reconstruct the detailed events that led to this incident. The aircraft was destroyed and the instrument rated private pilot was fatally injured. Visual meteorological conditions prevailed at the time of the accident. Winds were 170 degrees at 16 knots gusting to 22 knots. No flight plan was on file

> "The vertical and horizontal stabilizers had some skin wrinkling, but little evidence of ground impact. Both propellers displayed forward bending, chordwise dirt streaks and had dug into the ground, burying the spinners. No engine anomalies were found. No control anomalies were found. Fuel was present at the scene, and all tanks were ruptured. Fuel was found in the lines to both engines." [587]

This incident provides an extreme example of the uncertainty and ambiguity that can arise when investigators cannot access some of the sources of evidence mentioned in previous chapters. However, it is also important to stress that similar problems may also arise from the failure of data recorders. This topic was addressed in Chapter 7. Ambiguity is also likely to affect the initial stages of reconstruction before all of the available sources of evidence can be retrieved and analysed .

Ambiguity also occurs if there is genuine uncertainty about the events leading to an incident. For example, the following NTSB incident report describes how it may sometimes not be possible to resolve contradictions in witness statements:

> "During the takeoff roll directional control was lost and the aircraft rolled off the left side of the runway. Heavy braking was applied in order to stop short of a fence and the aircraft nosed over inverted. Both occupants were rated pilots. Their statements were contradictory. It was not determined which pilot was manipulating the controls or serving as pilot in charge at the time of the accident." [585]

In other circumstances, it is often possible to build a number of alternative reconstructions that reflect different hypotheses about the events leading to an incident. The apparent contradictions in the evidence can be addressed by constructing several models; each of which assumes that one particular version of events is the correct one. It is then possible to inspect the resulting reconstructions to determine which version of events is the most likely given the balance of evidence. For example, another NTSB incident report describes how a pilot lost control of their aircraft during an acrobatic maneuver [586]. Some witnesses stated that incident occurred when the aircraft was performing an outside loop. Others stated that the failure occurred during an inside loop. Two reconstructions can be developed to reflect each of these possible hypotheses about the sequence of events before this incident.

Ambiguity also arises when investigators cannot be confident in the evidence that they have obtained about the course of an incident. In extreme cases, this can arise when there are only third party statements about what might have happened. For instance, the following incident report relies upon a witness observation of an aircraft that has still not been located:

> "The pilot signed the pilot authorisation form to rent the airplane on December 25, 1994, about 13:25. Before departure both wing fuel tanks were filled at the request of the missing pilot. The time of departure has not been determined and there was no evidence of contact with any FAA ATC facility. A witness reported seeing a low wing airplane

> about 18:00 local 300-500 feet above ground level flying Westbound. He reported that the engine was sputtering when the airplane flew over his house. The missing airplane did not return to the departure airport..." [589]

In this case the narrative description that 'models' the course of events leading to the failure does not explicitly state that the aircraft observed by the witness was the missing Piper. This ambiguity is intentional; it may or may not have been this aircraft. It reflects the lack of certainty about the course of an incident whose causes could not be determined.

Ambiguity can be used to hide the underlying complexity of particular aspects of an incident. This is important if reconstructions are to provide investigators with an overview of an incident. If all of the details of a metallurgical or meteorological analysis were included then there is a danger that individuals might become 'bogged down' in less salient information. As a result, summaries are supported by further references to other documents that can be accessed to obtain additional detail if required. Chapter 14 will describe some of the problems that this style of reconstruction can cause for the readers of an incident report. For now it is sufficient to observe that ambiguity often occurs because of the abstraction or filtering process that is used to construct an overview of complex failures. As we shall see, however, it is critical that this process does not have the side-effect of hiding critical information about the course of events that contribute to a failure.

This section has presented a number of reasons why ambiguity can arise in the reconstruction of safety-critical incidents. In other words, we have shown that there are coherent reasons why investigators may simultaneously provide different accounts of the events that contribute to a single failure. However, this ambiguity also creates a number of potential problems. Firstly, there is a danger that any ambiguity in the initial stages of an investigation will not be adequately resolved by the time that a final report is issued. The previous paragraph described an incident in which it was not possible to identify the events that contributed to a aircraft going missing. In such circumstances, ambiguity cannot be adequately resolved and this is explicitly stated in the NTSB report. However, other incident reports are significantly weakened by ambiguities that seem to have been overlooked or ignored by the investigators. For instance, Johnson describes how one maritime incident report fails to describe what crew members were doing in the critical moments before a collision occurred [412]. In consequence many who read the report were left unconvinced about the investigators condemnations of the crews' actions during that interval.

There is also a danger that ambiguity can lead to misunderstanding. The use of ambiguity and abstraction supports several different interpretations of the meaning of a sentence. However, as a result there is a danger that investigators will read more into an account than was intended by the author. Conversely, they may fail to identify the intended meaning of a high-level reconstruction. It can be difficult to determine whether multiple interpretations reflect genuine uncertainty on the part of the writer or whether ambiguity is the result of necessary abstraction from underlying complexity. For example, some of the incident narratives cited in previous paragraphs do not provide information about meteorological conditions. Others omit information about the role of Air Traffic personnel. In the initial stages of an investigation, it can be difficult for the reader to know how to interpret these omissions. It might be assumed that there were no air traffic events contributed to the incident unless they are specifically mentioned. This interpretation need not be correct, for instance, if air traffic logs were still being assembled. Such problems can be minimised by introducing rules that force investigators to explicitly state when certain events did NOT contribute to an incident. For instance, NTSB incident synopses often exploit this approach. However, many regulators have introduced taxonomies that the categorise many different events that might lead to an adverse occurrence [717]. It is clearly impracticable to explicitly state when each of these events does not contribute to an incident.

A final problem is that ambiguity can arise from the medium in which a reconstruction is presented. As we have noted, natural language offers a flexible and expressive medium of communication. However, this power is achieved precisely because it permits ambiguity. Multiple interpretations are simultaneously supported by the use of imprecise language. Ideally this imprecision can be resolved in the final report on an incident by the introduction of additional evidence as it becomes available. However, as we have noted, there are many instances in which imprecision and ambiguity have persisted into the final versions of an incident report. There are further related problems. In particular, there are some properties that are inherently difficult to represent within natural language. For instance, it can be difficult to describe the way in which concurrent events can simultaneously occur across many of the different distributed systems that are involved in complex incidents. If these events are groups according to the systems that generated them then readers get a good idea of what happened to that particular system over time.

However, it can be difficult to gain an overview of what else was occurring throughout the application at any particular interval. Conversely, if a purely temporal sequence is exploited then it may be easier to see what events were happing at each moment in time. However, readers will have to piece together the individual events that occurred within a particular subsystem. A number of techniques can be used to address these limitations. For example, many incident reports contain text-based time-lines. These are constructed using a tabular form that lists the most salient events that contribute to a particular failure. These are recorded in the order in which they are presumed to have occurred during an incident. The investigator then notes down the time at which each event occurred as an entry in the table. Table 8.1.3 illustrates this approach. It provides an overview of the ways in which various events contributed to a derailment [603].

A limitation with this approach is that particular events can become lost amongst the many different items that are recorded in this tabular overview. As a result, some investigators also produce more detailed tabular time-lines that focus on the events that occurred in a particular subsystem or that influenced particular aspects of an incident. For example, Table 8.1.3 focuses on meteorological conditions during the grounding of a tug [601]: The time-line shown in Table 8.1.3 illustrates a similar approach [599]. In this incident, a table is used to chart the timing of an emergency response to a highway incident. There can, however, be considerable overheads involved in ensuring that these multiple time-lines provide a consistent account. It can also be difficult to ensure that any changes in the ordering of a particular time-line, such as those shown in tables 8.1.3 and 8.1.3, are reflected by consistent updates to an overall time-line, such as that shown in Table 8.1.3. These problems are compounded by the difficulty of gaining an accurate overview from many pages of prose descriptions. It can often be difficult to visualise the flow of events that contribute to particular adverse occurrences. What we need, therefore, are tools and techniques that can be used to explicitly capture properties of an incident, such as the temporal ordering of events, that are difficult to reconstruct using prose narratives. It should also be possible to use these descriptive techniques to identify any potential inconsistencies or ambiguities that might exist in the reconstruction. If these stem from a lack of evidence then either further investigations must be initiated or the final report must acknowledge the ambiguity. If inconsistencies are the result of clerical errors in drafting the report then they should be rectified. The following section introduces a range of techniques that have been proposed to satisfy these requirements for the reconstruction of incidents and accidents.

8.2 Types of Simulation

The previous pages have argued reconstruction techniques help to piece together the evidence that is derived from primary and secondary investigations. This, in turn, helps to determine whether necessary evidence is missing or whether there are any contradictions within the evidence that has already been gathered. Reconstruction techniques can also be used in a more subjunctive fashion. By this we mean that analysts can generalise beyond what is known about a particular incident to assess what might have happened under a number of alternative versions of events. As mentioned previously, we distinguish between two different forms of reconstruction. This chapter focuses on the use of computer-based simulations. These are an increasingly popular means of visualising the events that lead to safety-critical incidents. For instance, sketching tools can be used to derive simple story-boards. Alternatively, digital animation systems support more complex, interactive presentations of adverse occurrences. CAD-CAM tools can also be used to build detailed models of the behaviour of physical systems. Virtual reality systems can be used to drive immersive, interactive simulations with varying degrees of 'realism'. A distinguishing feature of all of these approaches is that they rely on computer-based tools to help analysts visualise the course of an incident. In contrast, Chapter 9 looks alternative approaches that do not rely so much on the use of computer-based simulations. A

Time	Events before the accident
01:43	BNSF receives flash flood warning (0001) for the Kingman area.
01:57	Track supervisor for Kingman area is notified.
02:24	National Weather Service issues severe thunderstorm and flash flood warning for central Mohave County, effective until 3:30 a.m. Also, before 3 a.m., weather updates (0002 and 0003) are issued to BNSF, including to watch for flash flooding, effective until 4:30 a.m.
03:39	Crew-members of westbound train Q-LACMEM1-08 report to the BNSF train dispatcher that the rain is letting up at Walapai (MP 501.3), and that they saw water in the culverts.
03:56	Train Q-LACMEM1-08 crewmembers at MP 489.7 report to the train dispatcher that there is no water on the ground and only trickles in the ditch.
04:05	The track supervisor begins his special inspection at MP 516.5, moving in an eastward direction.
04:12	Dispatcher tells track supervisor of Q-LACMEM1-08 information.
04:28	Contract weather service issues update 0004 to BNSF, advising to watch for flash flooding, until 6 a.m.
04:30-04:45	Track supervisor reports from Hackberry (MP 509.4) to the BNSF train dispatcher. He does not report high water. He inspects bridge 504.1. He notes water flowing adjacent to and under the bridge. He does not note any unusual track alignment or take exception to either the east- or westbound bridge.
05:07	Dispatcher reports to track supervisor that eastbound Amtrak train 4 is leaving Franconia (west of Kingman).
05:35	Westbound train B-CHCLAC1-05 passes Walapai (MP 501.3). Shortly thereafter, this train crosses the bridge on the north track at MP 504.1. Train crew notices nothing unusual about the bridge on the south track.
05:46	Track supervisor reports from Peach Springs (MP 465.8) to dispatcher. He says he will clear shortly for Amtrak train 4. He does not report any high water.
05:56	Amtrak train 4 derails at bridge 504.1S.

Table 8.2: Textual Time-line Reconstruction of Events Leading to a Derailment

1:57 p.m.:	winds S-SE/ 25 knots, seas 6 to 8 feet
about 3 p.m.:	winds 26 to 36 knots, seas 10 to 12 feet
4:30 p.m.:	seas 25 to 30 feet
5 p.m.:	winds S-SE /40 to 50 knots, seas 20 to 30 feet

Table 8.3: Textual Time-line of Meteorological Events in a Grounding

Time	Time from initial notification	Action
05:53	00:00	Initial 911 call received by WCSD dispatch
05:54	00:01	EMS dispatched
05:56	00:03	Two Slinger Police Department units arrived on scene
06:02	00:09	SFD command vehicle arrived on scene
06:07	00:14	First EMS unit arrived
06:19	00:26	Flight for Life dispatched from Milwaukee
06:38	00:45	Flight for Life arrived on scene
06:52	00:59	Ambulance delivered first van victim to area hospital
08:01	02:08	Flight for Life helicopter delivered second van victim to trauma center
12:28	06:35	Northbound lanes of US 41 reopened
14:03	08:10	Southbound lanes of US 41 reopened; area cleared

Table 8.4: Textual Time-line of Emergency Response to Road Accident

range of formal and semi-formal notations are used to model the events leading to a gas pipeline explosion. A number of further features distinguish these graphical and textual notations from the approaches in this chapter. In particular, the following computer-based simulations typically lack any formal underpinning. They are the product of iterative development and the subjective introspection of analysts. As we shall see, this provides great flexibility in the range of models that can be constructed. However, it also creates a number of practical problems. For example, in many virtual reality simulations it is entirely possible to break the rules of time and space. The same individual can be represented in two different places at the same time during an incident. In contrast, formal and semi-formal models are often supported by precise rules about what can, and what cannot, be represented. They provide mechanisms that help to identify omissions and inconsistencies, such as that mentioned above.

Some formal and semi-formal modelling techniques are supported by computer-based tools. Computer-based simulations can be directly derived from some abstract notations [719]. The distinctions between computer-based simulations and abstract models are, therefore, often blurred. These distinctions are retained, however, because they help to identify two different sets of concerns about the reconstruction of safety-critical incidents. For instance, formal mathematically-based notations can be difficult to interpret by non-mathematicians. In contrast, interactive simulations can often lead to unwarranted interpretations if analysts are seduced by particular animation techniques.

8.2.1 Declarative Simulations

One class of computer-based reconstructions can be described as ‘declarative simulations’. Declarative models describe aspects of a system that do not change during the course of an incident. At first sight, this definition seems to go the general idea that a reconstruction should provide an overview of the *events* leading to an incident. As we shall see, however, these declarative simulations can be used to illustrate the state of a system before and after an event. For instance, they can be used to illustrate the impact of a component failure upon the integrity of a hardware assembly. Sequences of these more static simulations can, therefore, be used to build up an impression of change over time. It is also possible to integrate declarative simulations with other forms of reconstruction to combine a static model of the system with more dynamic views of an incident.

Maps and Plans

Maps and plans provide important information about the environment in which many incidents take place. They can be annotated to denote the position of key objects and individuals before an incident occurs. Further annotations can be used to indicate the changing position of those objects at various moments during an incident. As with many of the techniques described in this chapter, it is perfectly possible to exploit this approach manually. We have, however, chosen to focus on computer-based techniques because they represent a significant area of innovation and development in incident reporting.

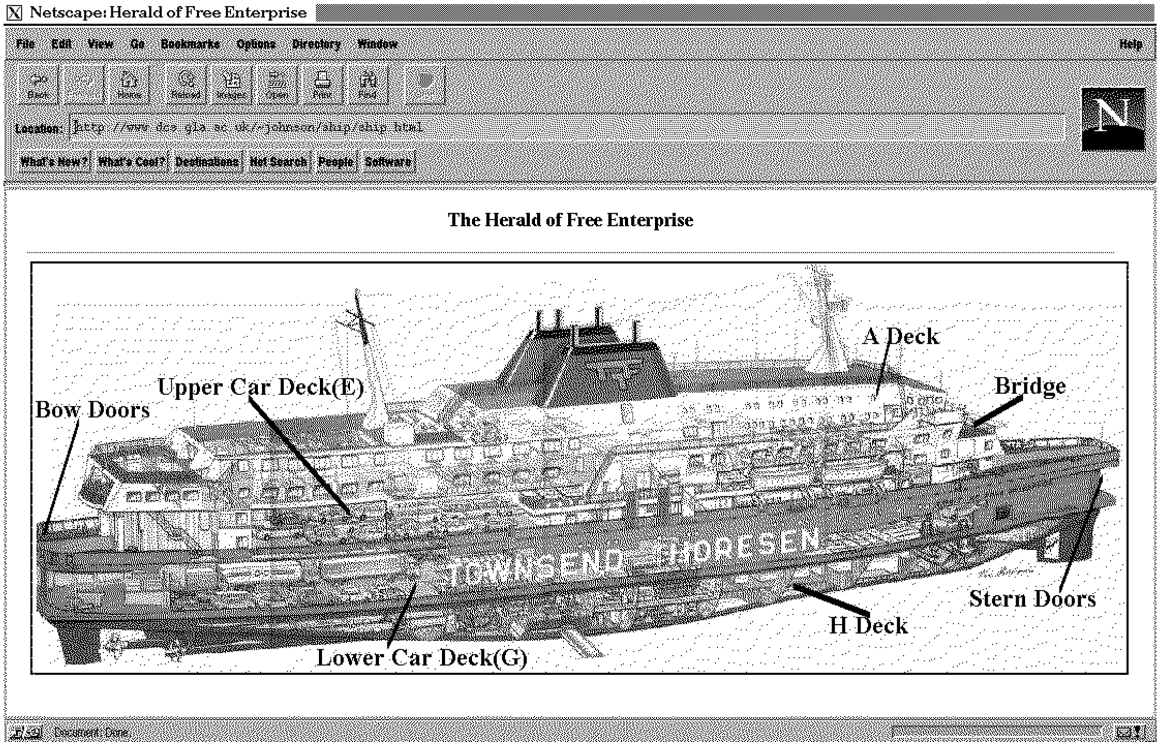

Netscape: Herald of Free Enterprise

Location: http://www.dcs.gla.ac.uk/~johnson/ship/ship.html

The Herald of Free Enterprise

Document: Done

Figure 8.1: Imagemap Overview of the Herald of Free Enterprise

Figure 8.1 illustrates a plan based approach to incident and accident reconstruction. It presents an imagemap of the Herald of Free Enterprise . Thanks are due to The Motor Ship and V. Berris (FSIAD) provided the sectional diagram of the Spirit of Free Enterprise. The interactive reconstruction was developed in collaboration with Anthony McGill [530]. The cross-sectional diagram provides an overview of the layout of the ship that would not have been possible from an external photograph. The labels are used to indicate areas of the ship that played a particularly significant role in the course of the accident. The image is presented in a web browser so that it is available to other investigators over an intranet. This technology provides additional advantages. For instance, users can select areas of the image to request more detailed information about particular areas of the ship.

Figure 8.2 shows how users can select particular areas of the vessel shown in Figure 8.1 to request more detailed information about the incident. This information can take a number of different forms.

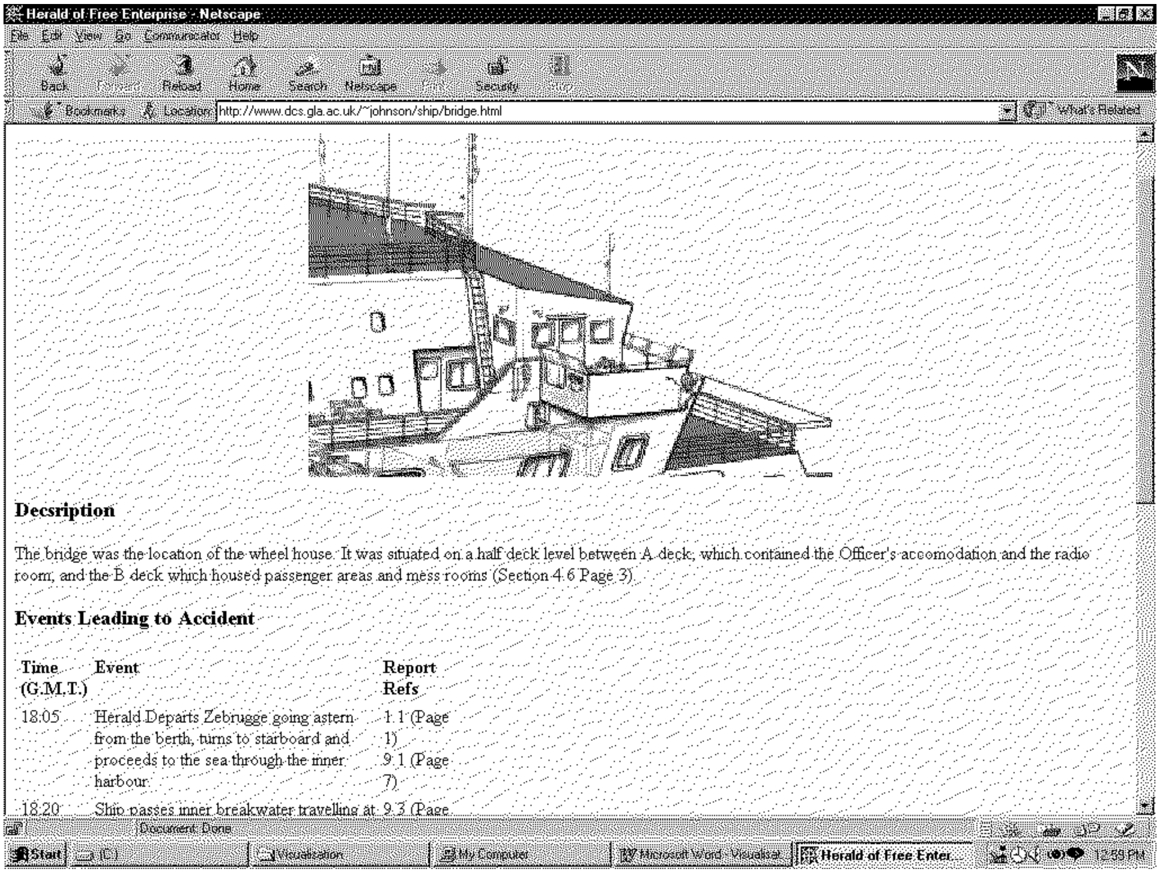
Herald of Free Enterprise - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://www.dcs.gla.ac.uk/~johnson/ship/bridge.html
What's Related
Document: Done
Start (C:) Visualisation My Computer Microsoft Word - Visualisat Herald of Free Enter... 12:59 PM

Decsription

The bridge was the location of the wheel house. It was situated on a half deck level between A deck, which contained the Officer's accomodation and the radio room, and the B deck which housed passenger areas and mess rooms (Section 4.6 Page 3).

Events Leading to Accident

Time (G.M.T.)	Event	Report Refs
18:05	Herald Departs Zebrugge going astern from the berth, turns to starboard and proceeds to the sea through the inner harbour.	1.1 (Page 1) 9.1 (Page 7)
18:20	Ship passes inner breakwater travelling at	9.3 (Page

Figure 8.2: Imagemap Detail of the Herald of Free Enterprise

For instance, if the user selects the bow doors they are presented with a range of engineering and construction information:

> "Situated at the bow of the G Deck, the bow doors were double weathertight doors of welded steel construction with a clear opening of 6.0m x 4.9m."

These static descriptions can also be augmented with information about the dynamic events that took place in a particular area of the ship. For example, if the user selects the bridge rather than the bow doors then the screen would be updated to show the textual time-line in Table 8.2.1. This illustrates the manner in which declarative plans, or maps, can be augmented with information about key events during the course of an incident. A time-line can be used to indicate when people and equipment move from one location to another. As can be seen from Table 8.2.1, events can also be annotated to indicate the evidence that supported each observation. This reconstruction was built post hoc and so the citations refer to paragraphs in the Sheen report [736]. However, investigators can exploit the same approach to keep track of the evidence provided by primary sources.

An important strength of the approach illustrated in Figures 8.1 and 8.2 is they avoid the 'god's eye' view that is provided by some computer-based reconstructions. Figure 8.2 only records information that was available to the crew on the bridge. It does not provide information that might have been available to crew on the car deck or in the passenger areas. In contrast, many other simulation techniques provide an overview of all of the evidence that is obtained during primary and secondary investigations. Such reconstructions provide a false impression because they integrate information that could not have been available to any single eye-witness. By using plans and maps to navigate into location-dependent time-lines, it is possible to gain a more accurate impression of

Time (G.M.T.)	Event	Report Refs
18:24	Captain sets combinator 6 on all three engines and the herald accelerates rapidly from 14 knots to possible ultimate speed of 18 knots.	9.3 (Page 7) IV (Page 71)
18:25	Steward hears water on the stairs.	IV (Page 71)
18:28	Ship lurches 30 degrees to port, temporarily becomes stable then slowly capsizes to port.	9.3 (Page 7) IV (Page 68) IV (Page 71)
18:28	Bridge clock stops.	IV (Page 71)

Table 8.5: Textual Time-line Integrated into A Declarative Reconstruction.

an individual's view of an incident. This was significant in the Herald of Free Enterprise accident because individuals on the bridge believed that another member of the crew was supervising the closure of the bow doors while he was actually asleep in his bunk. The 'god's eye' view can then be reconstructed by concatenating each of these discrete time-lines into a single sequence of events.

It is important to emphasise that these declarative models are important not simply in documenting the state of a system prior to an incident, they can also play an important role in documenting the consequences of particular failures. This is most apparent in the use of maps and plans to document key features of major accidents. The UK Air Accident Investigation Branch provide an innovative example of this approach, accessible via [9]. They used a computer-aided design system to model the hull of a Boeing 747. Sections of this model were then annotated to denote the locations where investigators found components on the ground around Lockerbie.

There are more complex examples of map based techniques being used during the reconstruction of adverse occurrences. For instance, police investigators must survey the markings that vehicles leave on a carriageway following an incident [442]. Traditionally, this has involved the use of pencil and paper. Increasingly, however, digital equipment is being used to capture the position of those markings in relation to the layout of the road. This information is then directly downloaded into computer-based reconstruction software. A key aspect of this approach is that many systems support *backwards reasoning*. Information about pre-incident events can be deduced from observations about the consequences of an incident. Many systems use skid marks to deduce the speed and trajectory of particular vehicles. There are, however, a number of limitations with the general application of this approach. Forwards and backwards reasoning from declarative reconstructions can introduce uncertainty. For instance, if the skid marks are eroded or obscured by the effects of the weather or by other debris then it may not be possible to have complete confidence in the results of any consequence calculations. At a more basic level, the calculations that are used to calculate speed and velocity from road markings are, typically, governed by confidence levels. As we shall see, this caveat has a number of important consequences. For example, animated reconstructions can be derived from the information that is deduced by map-based surveying systems. These animations present vehicles travelling at a particular simulated velocities that do not reflect the degree of uncertainty in the initial calculations.

There are further problems with the plan or map based approaches described in this section. Previous figures used an image of the Spirit of Free Enterprise. This is the sister-ship to the vessel that was actually involved in the capsize. We were compelled to use this image in our reconstruction because there is no similar image available for the Herald of Free Enterprise. The sister ship was also used by the official court of enquiry. This example illustrates a more general point. In the aftermath of an incident it may not be possible to obtain a detailed plan or map of particular locations. It may be possible to produce a sketch of the probable location of key objects and people. However, it can be difficult to represent the fact that such locations are based around inferences rather than direct evidence about the scene of an incident. Similarly, the process of developing cross-sectional sketches, such as that shown in Figure 8.1 can introduce biases and distortions of perspective. Photorealistic simulations reduce some of these problems.

Photorealistic Models

s mentioned, plans and maps have long been used to help analysts reconstruct the location in which an incident occurred. Artists sketches can, however, provide a poor impression of the environment, objects and individuals who contribute to safety problems. As a result, photographs are often used to supplement maps and plans. There images provide a more direct impression of the location in which an incident occurred. They can be used to provide detailed information about the physical state of process components, whether they are new or worn, whether they are correctly installed or misaligned, whether they were damaged by an incident or whether they remain intact. Photographic images can also provide an impression of particular environmental factors, such as the line of sight between an operators and a warning signal. Such information can be difficult to convey using plans and maps. It is important to emphasise that such techniques are not immune from some of the biases that sketched images. Different camera angles, exposures and processing techniques can give false impressions about what could or could not be observed during the course of an incident.

As with plan-based simulations, there has been a recent revolution in the use of photographical images to support incident reconstruction. A range of computer-based techniques are creating both opportunities and challenges for investigators. One of the biggest benefits of recent developments is that investigators can take images in the field using digital cameras. These can then be sent from a laptop PC using a modem and wireless telephony to colleagues in other regions country and throughout the globe. in some recent cases this has been done interactively with the field investigator being guided remotely to take live images of the incident site. The resulting photographs can then be archived on servers that can then be accessed by other investigators when they are needed. These images can also be used in subsequent litigation and in any subsequent reports.

As mentioned, these benefits also bring a certain number of concerns. It is possible to falsify conventional photographic images but digital editing techniques make this far easier. A vast range of 'post-production' effects can be achieved with relatively little training. In consequence, many countries have strict rules about the ways in which digital resources can be used both during an incident investigation and during any subsequent litigation. For example, investigators can be required to testify that digital resources have been protected from unauthorised 'tampering'. Digital watermark techniques provide one means of achieving this protection. These watermarks are implemented using an identification code that is permanently embedded into electronic data. The code carries information about copyright protection and data authentication. If the resource is edited in any way then the watermark will be destroyed. Evidence of 'tampering' is then apparent if users cannot extract the original watermark from the electronic resource. The interested reader is directed to Hartung and Kutter's overview of multimedia protection techniques for more information about this and similar approaches [310]. In contrast, the remainder of this section focuses on the use of photorealistic pseudo-3D techniques for incident reconstruction.

One of the problems with static images is that it can be difficult to gain an impression of three dimensional space. This is a limitation of both conventional and digital images. Analysts cannot use these images to 'walk around' the scene of an incident. Instead, they are restricted to the perspectives chosen by the person taking the photograph. This becomes an issue because it can be difficult for investigators to predict all of the possible perspectives that might be relevant in the aftermath of an incident. Even if they expend vast amounts of time and film, it can still be difficult for other investigators to piece together the layout of an environment from dozens of static images and plans. Software engineers have responded to similar concerns within other application areas by developing photorealistic tools for desktop virtual reality (desktopVR). These tools avoid the use of cumbersome helmets and gloves that have been recruited by immersive virtual reality systems. Instead, they attempt to provide an impression of movement in 3D space using conventional input and output devices; keyboards, mice and standard computer displays. Although we focus on this non-immersive approach to virtual reality, many of the comments also apply to incident reconstruction systems that expect their users to wear helmets, gloves and other more complex apparatus. In contrast, QuicktimeVR constructs an interactive simulation from a number digital images and presents them

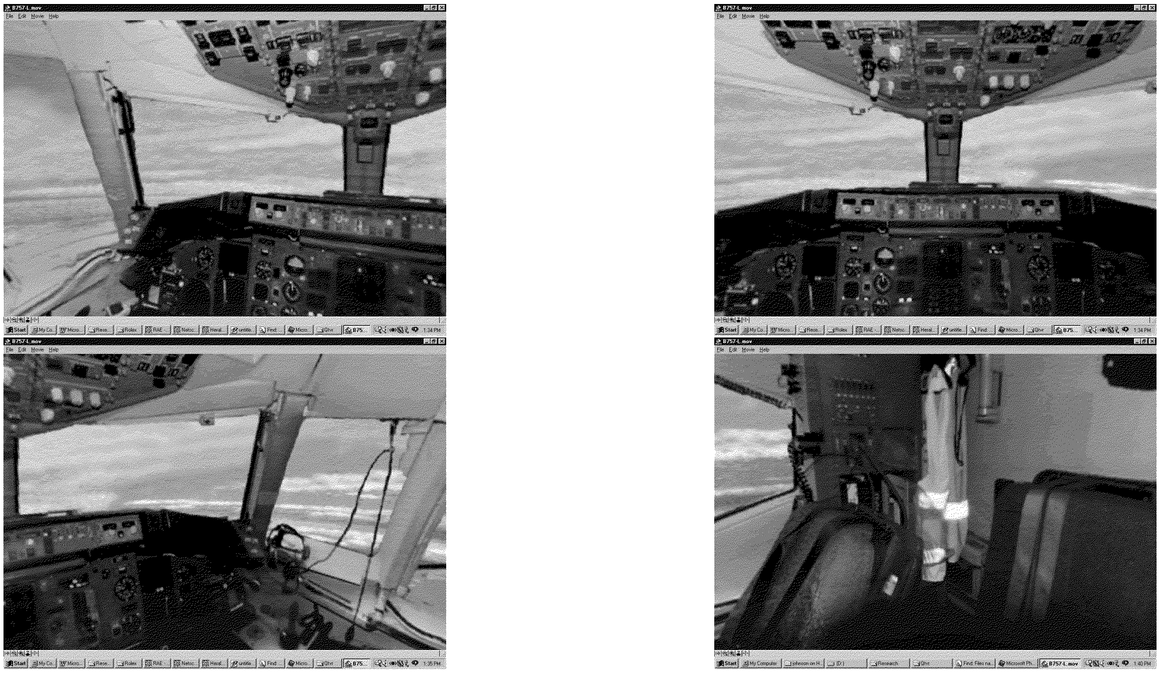

Figure 8.3: QuicktimeVR Simulation of a Boeing 757

on standard desktop displays. These photographs are taken using a motorized tripod which ensures that a still image is taken approximately every N degrees. The value of N is determined by a number of factors including the type of lens that is being used as well as the distance from the camera to the visual horizon. These photographs are then 'stitched' together by the interpolation software. The net effect is to enable users to pan through 360 degrees simply by holding their mouse over the image. Digital effects can also be introduced so that users can zoom in to view particular details of each image. The idea of motion is provided by moving the tripod to a different location. The process described above is then repeated so that the user can again pan around to view the environment from the new location. Figure 8.3 presents images taken from a QuicktimeVR simulation of a Boeing aircraft [425]. The QuicktimeVR images in this section have been reproduced with kind permission of Strathclyde Regional Fire Brigade and were produced in collaboration with Bryan Mathers, Alan Thompson and Bill West [524]. It should be noted that these images provide an extremely poor impression of what is like to interact with the desktopVR system. The resulting simulations not only help incident investigation. They can also provide lawyers, jurors and engineers with an impression of hazardous environments. Unlike films and videos the interactive nature of this approach also enables users to choose their own path through the scene of an incident.

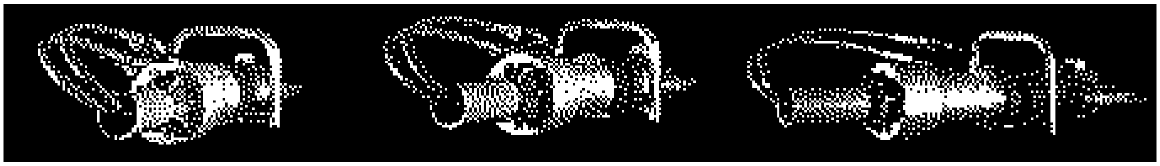

Figure 8.4: QuicktimeVR Simulation of Lukas Spreaders

Figure 8.3 illustrated how the application of desktopVR techniques can be extended from the domains of computer aided learning and scientific visualisation to support incident reconstruction. The Boeing cockpit illustrated in these images represents one of two particular approaches to the QuicktimeVR technique. In this example, a tripod is moved around taking images that form 360 degree sweeps on the scene of an incident. The same approach can, however, also be used with slight variations to provide interactive simulations of process components. This is illustrated by Figure 8.4. Digital editing techniques have been used to joint together a number of still images from a QuicktimeVR reconstruction. The intention has been to provide an impression of the way in which a user can manipulate the software to rotate the object through 360 degrees. The Figure shows how this approach has been applied to the Lukas spreaders that Fire Crews use to extract passengers from road traffic accidents. The process used to create these simulations is slightly different than that used to produce the system in Figure 8.3. Rather than rotating a camera to photograph the environment, in this approach the camera is typically held in a constant position while the object is rotated. By raising or lowering the position of the camera, it is possible for the user to rotate the object around both the x and the y-axes. They can zoom in to view both the top and the bottom of the object.

These photorealistic simulation techniques are declarative because they provide a pseudo-3D impression of the state of the environment or of critical objects at a particular instant in time. However, it is possible to construct multiple simulations to show the state of an object or environment both before and after key events. We have used this most frequently to show the effects of damage or wear on process components. It is also important to emphasise that these resources can easily be integrated into other electronic resources. They can be 'marked-up' in the same way that we annotated the sketches of the Herald of Free Enterprise. This enables users to electronically access linked pages of textual information by selecting areas of the images. Investigators can use this approach to access to the text-based time-lines illustrated in the previous section. This provides an alternative means of introducing dynamic information about critical events into what would otherwise be a static and declarative approach. There are further benefits. For instance, we have used heavy-duty turntables to represent artifacts that are too heavy to be lifted. This enables users to inspect objects such as aircraft engines. They can literally turn these component 'upside down' to select the best angle from which to view any potential damage. Similarly, we are exploring the use of miniaturised cameras to create pseudo-3D models of devices that cannot normally be directly accessed.

At the time of writing this book, these techniques have not been widely exploited to support incident reconstruction. It seems likely that this will change. We are using techniques that are being developed for mass-market computer aided learning systems rather than complex bespoke techniques for incident reconstruction. As a result, these simulations can be developed using relatively low cost technology. QuicktimeVR simulations can be produced at a fraction of the cost of a traditional sketched plan. The examples illustrated in this section were produced by two undergraduate students working in collaboration with members of Strathclyde Fire Brigade. Members of this service have since gone on to develop their own applications of this technology. It is also important to note that the costs of producing these simulations can be defrayed by their multiple applications. All of the systems illustrated in this section have been integrated into the Fire Service's computer-based training schemes.

8.2.2 Animated Simulations

The previous section has presented a range of techniques that support the reconstruction of safety-critical incidents. Maps and plans provide represent the layout of an environment and can be used to locate objects within it. Photorealistic techniques augment these sketches and provide a richer source of information about the scene of an incident. For instance, they can provide evidence about the location of objects that might otherwise have been overlooked when a map is being sketched. However, both of these approaches provide a declarative snapshot of the context in which an incident occurred. They rely upon secondary techniques, such as textual time-lines, to represent more dynamic information. In contrast, the simulation techniques in this section are specifically intended to help analysts reconstruct the changing events that contribute to safety-critical incidents.

Physical Simulations

Previous sections have mentioned that many incident investigations rely upon full scale reconstructions of adverse occurrences. The scale and sophistication of such reconstructions varies from impromptu demonstrations by investigators in the field to full-scale simulations using highly expensive facilities, such as the US Department of Labour's Mine Simulation Laboratory. This facility provides training for safety inspectors and incident investigators using a 48,000 square foot above-ground simulated simulated coal mine. Similarly, NASA's Langley and Ames research centres operate a number of aircraft simulators that attempt to provide a pilot with enough sensory information to convince the pilot that an actual aircraft is being flown. Sensory cues include realistic out-the-window scene generation with 360 degree field of view. The more advanced facilities also provide motion feedback with accelerations applied to the cockpit to simulate momentum shifts in an aircraft. They also employ special seats and suits that are intended to mimic gravitational pull. Other forms of tactile information is simulated by positive force feedback in the pilot's stick and pedals. Acoustic systems simulate natural sounds, such as the wind, and artifical sounds including realistic engine profiles,

Physical reconstructions often involve a high-degree of risk. Reconstructing failures that almost led to a disaster can lead to disaster. Physical reconstructions also, typically, associated with high costs. The NTSB's Wake Vortex flight tests provide an extreme example of this. These were conducted at the FAA's Technical Center in Atlantic City, New Jersey in September, 1995. These vortices can be thought of as a form of turbulence. They can be created whenever an aircraft passes through a section of airspace:

> "Wake vortex: A counterrotating airmass trailing from an airplanes wing tips. The strength of the vortex is governed by the weight, speed, and shape of the wing of the generating aircraft; the greatest strength occurs when the wings of the generating aircraft are producing the most lift, that is, when the aircraft is heavy, in a clean configuration, and at a slow airspeed. (Also known as wake turbulence.)" [609]

Air traffic controllers must, therefore, follow wake vortex regulations that specify minimum separations between particular types of aircraft. These regulations are intended to prevent the following plane from flying through a vortex created by its predecessor [367]. The NTSB's physical reconstructions used a Boeing 727 to generate a vortex. A Boeing 737 was then deliberately flown into the 727's vortex. These were identified using wing-mounted smoke generators on the lead aircraft. The results of the reconstruction were monitored by on-board sensors and from a T-33 chase plane. The risks of such simulations are obvious; especially considering that they were triggered by hypotheses about the causes of a number of previous incidents [609].

The NTSB's wake vortex studies typify the way in which physical simulations are used to obtain many different forms of data. Information was obtained from videotapes, an enhanced flight data recorder, from Boeings portable airborne digital data system, for a 2-hour cockpit voice recorder on the 737 and from test pilot statements. These data revealed that the 727 wake vortices remained intact as much as 6 to 8 miles behind the wake-generating airplane, and wake strength values ranged from 800 to 1,500 feet per second. The video tapes revealed numerous examples of wake vortices breaking apart; linking up; and moving up, down and sideways. This study is also remarkable for the way in which the NTSB exploited the video recordings. These were used not simply to support the investigators' analytical work. They were also used to provide the public with important insights into the conclusions of the final report. This seems to be an increasing trend as incident and accident investigators become increasingly aware of the need to communicate their findings beyond the regulators of their industry [749].

The principle aim of recording this information is to prove or disprove hypotheses about the causes of an incident. The NTSB study was conducted because wake vortices had been implicated in a number of previous incidents. This reiterates the close links between reconstruction and causal analysis. Physical reconstructions help to prove or disprove causal hypotheses. Unfortunately, however, the results from these studies can be highly ambiguous. The failure to recreate an incident may be due to characteristics of the simulation environment rather than weaknesses in the underlying hypotheses. There may also be problems in instrumenting a physical reconstruction to determine whether or not an incident has actually be recreated. Typically, the components and systems that were involved in an incident are heavily instrumented to provide as mush feedback about the potential causes of a failure as possible. This is critical because of the risk and expense mentioned above. Lechowicz and Hunt provide an example of this instrumentation when they introduce a system to provides in-motion weighing, load distribution analysis plus defect detection and classification at wheel, bogie, wagon and train levels [483]. In order to calibrate the measurements from their system they had to run a range of satisfactory and specially assembled bogies over the track. These assemblies provided a mixture of good wheels, defective wheels and vehicle loads and were run at speeds from 30 up to 130km/h. Such calibration tests must be conducted *before* an incident can even be simulated. If they are not then there is a danger that incorrect conclusions will be drawn from the results that are provided by the instrumentation.

Physical simulations supplement any information that can be derived from data recorders that were running at the time of an incident. This live 'incident' data also helps to validate the information that is obtained from physical simulations. If the information derived from a physical simulation is not consistent with that gathered at the time of the incident then a number of potential problems might be diagnosed. For instance, the physical simulation may not have recreated the conditions that held at the time of the incident. Alternatively, the recording devices that were operating at the time of the incident may have been incorrectly calibrated. The recording devices that were running during the incident may not have been correctly calibrated. However, such inconsistencies are not without value. They can often be used to identify alternative conditions that might lead to similar consequences during an incident.

The data that is derived from physical simulations and from incident data recorders can provide parameters for computer-generated models. This has significant advantages. For example, the costs associated with crashing rail trucks at different speeds, typically, prevents a wide range of physical simulations. However, these studies can be used as data points for computer-based simulations that can be run and re-run without significant additional costs. Later sections will identify the strengths and weaknesses of these techniques. In contrast, the following section focuses on the animated presentations that are frequently derived from these computer-based models. In many cases these models are used without the support of physical reconstructions. The high costs, the inherent risks and the problems of instrumenting a reconstruction often persuade investigators to rely upon solely computer-generated simulations.

Computer-Based Animations

We have shown how desktopVR software can 'stitch' together photographic images in order to provide interactive tours of particular environments or rotational images of process components. This is not the only way in which desktopVR technology might support incident reconstruction. In particular, model-based approaches can be used to reconstruct environments that cannot be directly photographed. This is often the case when an incident has resulted in significant damage to a workplace or if that environment has become hazardous in the aftermath of an incident. Model-based approach construct complex process components from a number of geometric primitives, such as spheres and cubes. Application software then renders the image of these composite objects onto the user's screen. The image is updated as the investigator moves around in the virtual environment. In particular, the rendering software must cope with changes in perspective and the occlusion that occurs when one object obscures the image of another.

Figure 8.5 illustrates how this model-based approach can be used to support incident reconstruction. It shows how geometric primitives can be combined to reconstruct complex structures including vehicles and people as well as buildings. This Virtual Reality Markup Language (VRML) simulation was developed in collaboration with Marcus Kramer [461]. This application enable analysts to place construction equipment within a number of different layouts that were proposed for a building site. The equipment could be moved around within these layouts to show users the potential hazards that were posed, for example by overhead cables.

Many model-based reconstructions, such as the one shown above, are developed using declarative environments. They enable users to reconstruct the topography of environment. Libraries can also

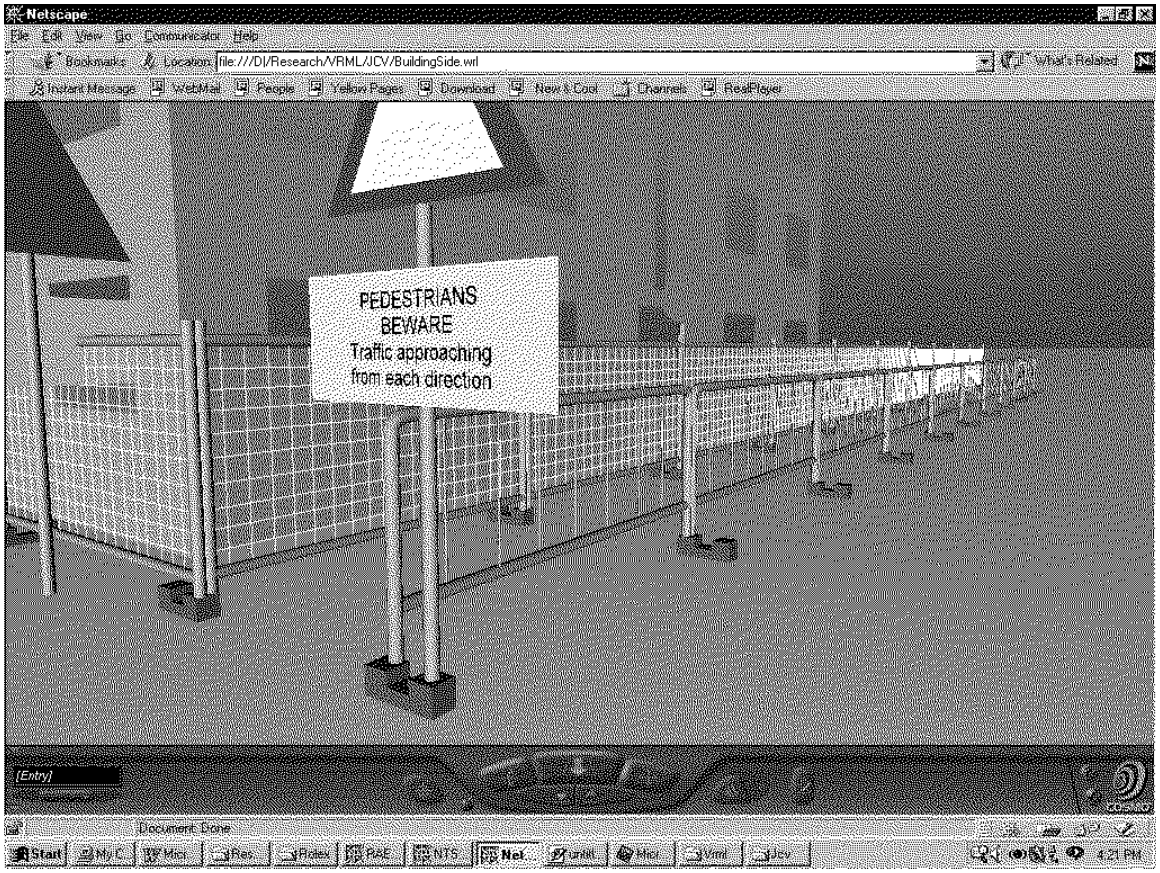

Location: file:///DI/Research/VRML/JCV/BuildingSide.wrl

Figure 8.5: VRML Simulation of Building Site Incidents

be constructed to provide easy access to a range of common process components that have previously been developed using the geometric primitives. Analysts must use scripting techniques if these static models are to recreate the 'real-time' behaviour of physical objects. Typically, these programming languages are event driven. Programmers specify the actions that are to be taken when the state of the model changes. For instance, it might be specified that the walls of a building in Figure 8.5 are deformed when a piece of construction equipment hits is driven into it. These programs are developed into simulations in an iterative manner. Any problems in the reconstruction are detected when the program is run and the simulation is viewed. The program can then be amended before being the model is viewed again. Ultimately, however, investigators derive digital animations which simulate the probable course of events leading to an incident. Figure 8.6 presents a number of still images from an NTSB reconstruction that used this approach [610]. This reconstruction had a dual purpose. It was used to validate the inspectors' view of the incident. It was also used to communicate their view of the incident at a public hearing that was convened by the NTSB following the investigation.

A number of concerns affect the production and use of digitised videos from model-based incident simulations. The examples in this section have been produced using extremely flexible modelling software. It is possible to position the 'camera' that determines the user's view of an incident at almost any point in three-dimensional space. This flexibility can result in considerable usability problems for incident investigators who are unlikely to be experts in the manipulation of these cameras. Small changes to the positioning of the viewpoint can result in users entirely missing key aspects of a simulation. As a result, the examples cited in this section use simulations to produce digitised videos that illustrate particular points about the course of an incident. This implies that

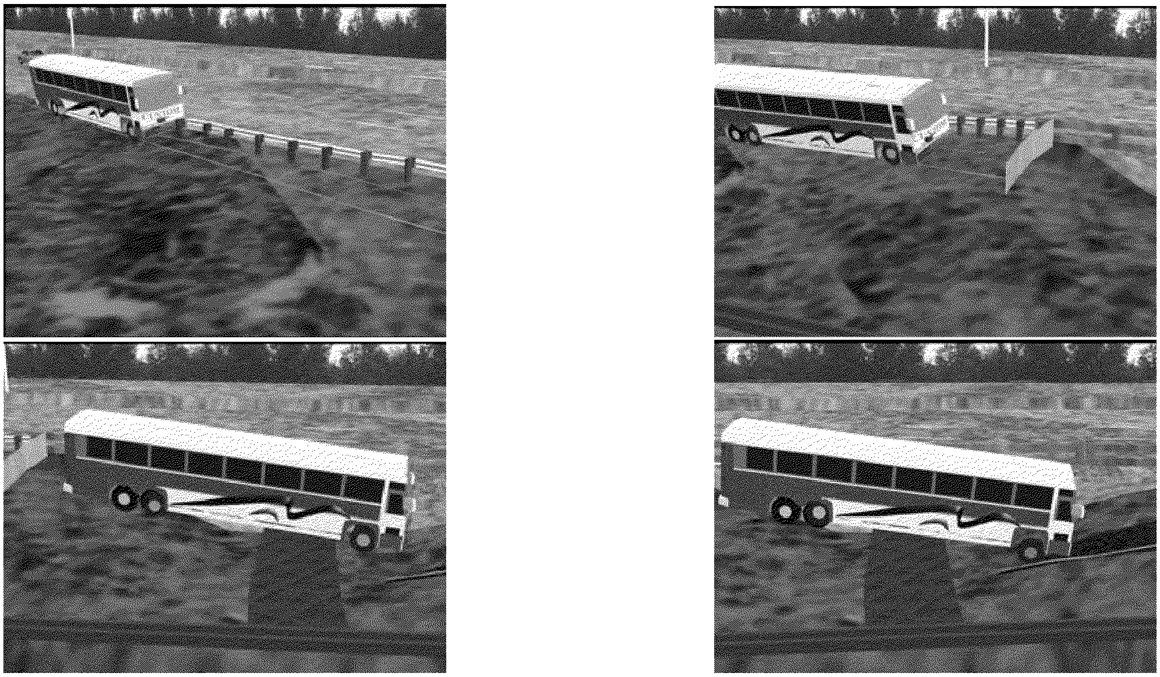

Figure 8.6: NTSB Simulation of the Bus Accident (HWY-99-M-H017).

the pictures obtained from the model are carefully edited to form a linear sequences of images. The end-user cannot change the camera angle nor can then move within a model to obtain new perspectives on an incident. This introduces concerns that investigators may produce videos that are biased towards particular aspects of a simulation. The viewpoint may be place close to the ground to emphasise the apparent speed of a vehicle. Alternatively, the camera may be 'locked' at the same velocity as a moving object so that the viewer gets less of an impression of the relative speed of that object.

Animated models also suffer from many of the concerns that were raised about digitised photographs. It is possible to create a false impression of the events and conditions that contributed to an incident. For instance, model-based reconstructions rely entirely upon the developer to specify lighting parameters. Programmer must ensure that there is a clear correspondence between the world that they *build* and the real situation in which the incident took place. It is for this reason that the NTSB presentation shown in Figure 8.6 is accompanied by the following notice:

> "Disclaimer: Simulations presented below used scene and geological surveys, highway design plans, witness statements, vehicle testing, vehicle plans and vehicle operating characteristics. The depictions represent actual lighting and weather conditions at the time of the accident." [610]

The need to ensure a correspondence between virtual components and physical objects has led investigation agencies and software developers to introduce numerical simulation routines into the scripting languages that drive model-based simulations. The NTSB's models were constructed using tools that were developed by the Engineering Dynamics Corporation. There tools enabled investigators to specify that the bus was initially travelling at sixty miles per hour. The simulation software then reflects the way in which the speed of the bus increased to sixty-two miles per hour when it hit the guardrail. The vault speed of the bus was set at fifty six miles per hour and the impact speed with the opposite side embankment was simulated fifty seven miles per hour.

In spite of the fact that many of these techniques are being applied by the NTSB and other investigation authorities, there are a number of important limitations with this approach. There are considerable costs associated with the time that is required to model even relatively simple objects in model based virtual environments. These costs can be reduced by maintaining libraries of common components. The elements in such libraries may not reflect the subtle differences that exist, for instance between particular models or types of equipment. These costs can dissuade analysts from reconstructing distal causes of adverse occurrences. DesktopVR reconstructions are often biased towards the simulation of those proximal events that have the greatest impact upon their viewer. It can also be extremely difficult to build interactive model-based simulations for certain classes of events. For instance, it is rare to see such simulations that show how managerial or regulatory decisions have influenced the course of an incident. Further problems relate to the use of model-based simulations to represent the outcome of an incident. The NTSB investigators had several good reasons when they decided not to model the impact damage and motion to final rest in Figure 8.6. Firstly, there are computational and technical difficulties in calculating the many different forces that act on complex objects during safety-critical incidents. As a result, simulations must approximate the mechanical and kinematic forces that operate on physical components. This is sufficient for most purposes but can lead to 'unrealistic' effects during impact sequences. More importantly, there are also strong ethical concerns that make the simulation of such consequences unacceptible to those who are involved in an incident.

Abstract Visualisations of Critical Events

Figure 8.7 presents a less familiar application of desktopVR for incident simulation. The interactive reconstruction was developed in collaboration with Anthony McGill [530]. Instead of modelling the objects that are involved in an incident, this approach provides a more abstract overview of the events leading to an adverse occurrence. The image shows three time-lines that recede into the z-plane. Each time-line describes events that are related to a particular aspect of the incident. In this case, systems engineering failures are distinguished from the actions of the chief officer and the assistant bosun. As can be seen, geometric primitives are used to model flags. These are labelled with information about events that affected the course of the incident. Each flag is placed at a point on the time-line. Users can exploit the application software to 'walk' forward along either of the lines. By looking horizontally across the x-plane it is possible to review those events that concurrently affect the other entities that are represented by the parallel time-lines. By looking forwards into the z-plane, the user reviews those events that happen in the immediate future from their position on the line. By turning 180 degrees in the same plane, they can review those events that happened immediately before their current position in time.

As can be seen from Figure 8.7, this technology can be integrated within conventional web browsers. As a result, hypertext links can be associated with the individual objects in the model. Users can select any of the flags to view a page of information about the evidence that relates to that event. This is essentially the same mechanism that was described for the QuicktimeVR visualisations introduced in previous sections. A number of refinements have been proposed for the basic approach described above. For example, we have set some of the flags at a 45 degree angle to the view shown in Figure 8.7. This has been used to provide viewers with a means of identifying when there is conflicting evidence about a particular event. Similarly, some flags are not planted into a time-line. Instead, they 'hover' above the rest of the flags. This has been used to denote events for which there is not accurate timing information.

The approach shown in Figure 8.8 has a number of important benefits. The three-dimensional structures used in the model are extremely simple. Analysts, therefore, avoid the overheads associated with the three-dimensional modelling of complex real-world objects illustrated by Figure 8.5. It is a trivial exercise to add new flags. Investigators simply duplicate the existing components and update its label. There are, however, a number of practical limitations. For instance, it can be a non-trivial exercise to use two-dimensional input devices, such as the conventional keyboard and mouse, to navigate three-dimensional space. These difficulties can be eased by the provision of a number of simple interface design techniques. For example, pre-defined viewpoints can be built into the system so that users can move from one point to another by selecting from a menu of options. This avoids the need for investigators to continually check whether any mouse movements will alter their position or orientation in the X,Y or Z planes. In Figure 8.7 this is achieved by a menu

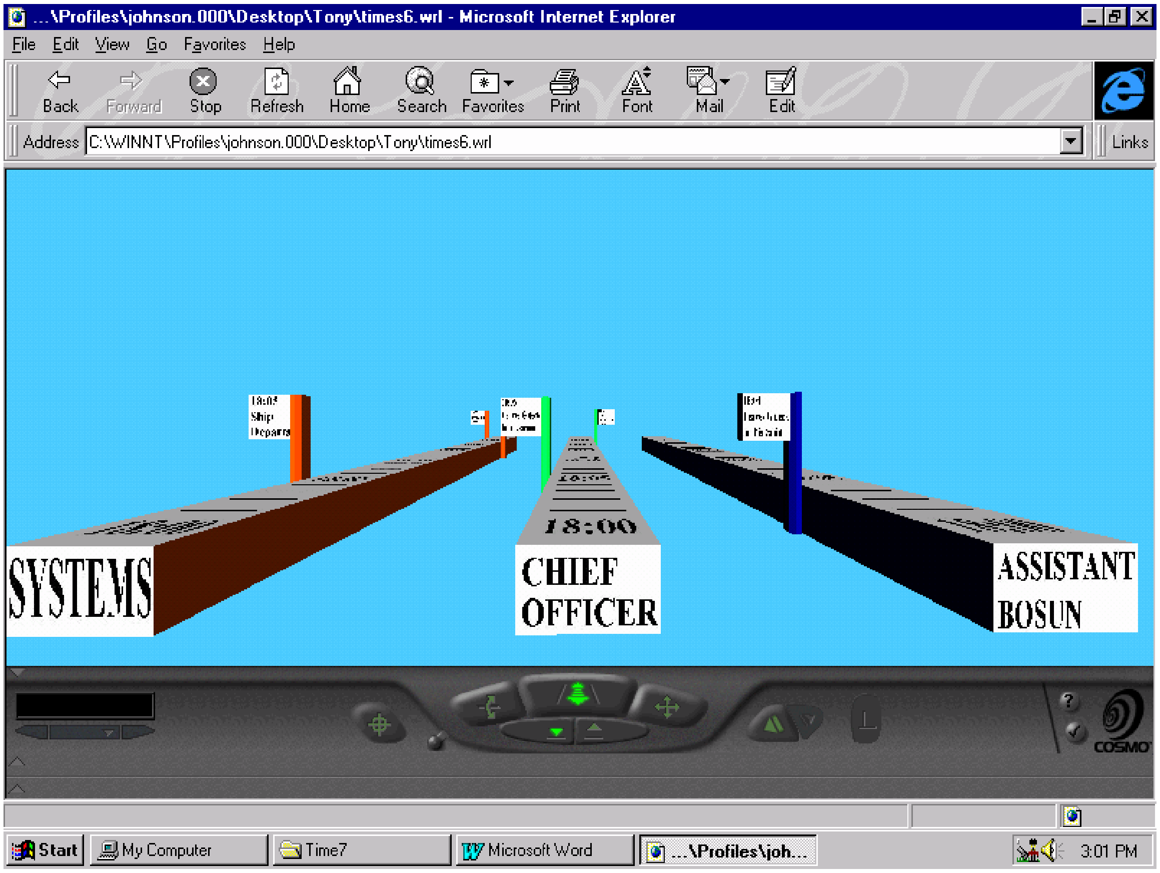

Figure 8.7: 3 Dimensional Time-line Using DesktopVR

of times. If the user selects 18:45 then the system automatically ‘walks’ them along the time-line to that position. Such techniques can be used to address specific navigation problems. They are not, however, a panacea and it is important to balance the enthusiasm for these new approaches against these practical problems. We cannot expect current generations of incident investigators to instantaneously acquire the three-dimensional skills of computer games enthusiasts.

The techniques illustrated in Figure 8.7 remain the subject of on-going research. Further applications must be developed to demonstrate that they can reconstruct a range of incidents in different industries. Further studies are also needed to determine whether or not there are alternative, more appropriate visualisation techniques. For instance, Figure 8.8 shows how Mackinlay, Robertson and Card’s idea of a ‘perspective wall’ can be applied to support incident reconstruction [509]. This model was developed in collaboration with Ariane Herbulot from the Ecole Superieure en Sciences Informatiques, Sophia Antipolis. The perspective wall makes greater use of the y-plane than the 3-D time-line illustrated in Figure 8.7. However, many of the underlying principles remain the same. Three time-lines are represented on the lower surface of the wall. Events are placed on one of the three parallel lines to denote whether they are associated with a particular subsystem or individual. Their position on the line denotes the moment at which they are assumed to occur. The upper wall is annotated with ‘real-time’ markers that act as reference points.

Figure 8.9 presents a detail from the perspective wall shown in Figure 8.8. As before, the visualisation is integrated into a web browser. Pages of additional information and evidence can be accessed by selecting individual events. This approach provides a number of benefits over 3-D time-line. As can be seen, concurrent events appear immediately above each other in the x-plane. They can, therefore, be seen at a glance from the position shown in Figure 8.9, which appears in

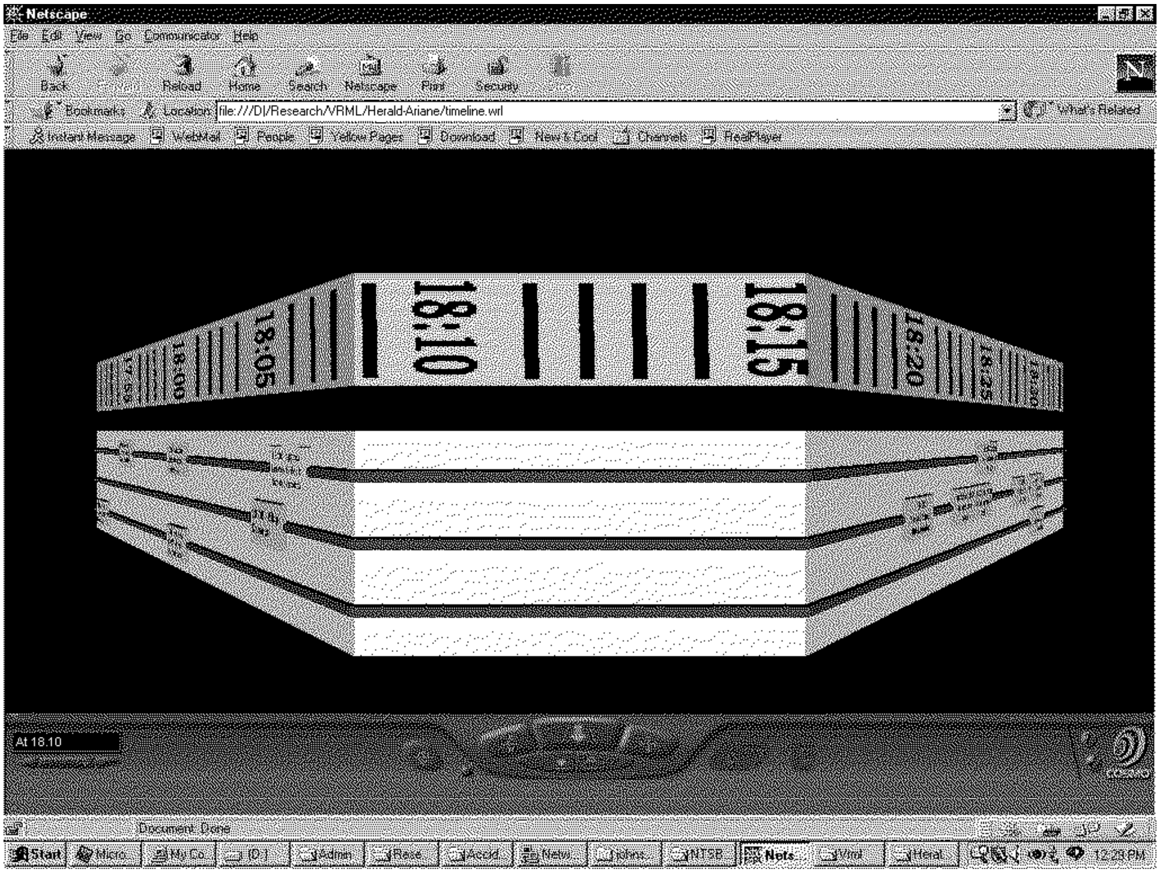
Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: file:///D|/Research/VRML/Herald-Ariane/timeline.wrl
What's Related
Instant Message WebMail People Yellow Pages Download New & Cool Channels RealPlayer
18:00
18:05
18:10
18:15
18:20
18:25
At 18.10
Document: Done
Start
12:29 PM

Figure 8.8: Overview of Perspective Wall Using DesktopVR

the menu of viewpoints that was described earlier. In Figure 8.7 concurrent events appeared on the same plance in the z-axis. In consequence, users would have to perform a more complex rotation in order to view these markers. This problem can be avoided if the view point is altered. For example, the analyst can easily see concurrent events on the 3-D time-line if they are positioned directly above the three time-lines. Experience has, however, shown that this can have a disorienting effect on the viewer.

Previous sections have argued that the application of low-cost photographic reconstruction techniques, such as QuicktimeVR, can be extended from mass-market publications to support engineering applications in incident reconstruction. This section has shown how techniques from scientific visualisation and information science, such as the perspective wall, can be applied to similar ends. The meta-level point is that these techniques look beyond current approaches to incident reconstruction. Many of these 'traditional' approaches owe more to the nineteenth century than to the twenty-first. This is a significant and forceful argument because it builds on the underlying analysis of Perrow [675] and Sagan [718]. Chapter 1 summarises their argument that the increasing complexity of many modern systems is leading to increasingly complex failure modes. It is certain that many reconstructions stretch the limits of what can be achieved using manual techniques. The analysis of the Allentown incident, discussed in Chapter 9, identified more than 1000 events contributing to human 'error', systems 'failure' and managerial 'weakness'. Some of these events stretched back more than five years before the explosion that triggered the incident investigation. It seems likely that computer-based visualisations will be necessary if investigators are to cope with the burdens imposed by the reconstruction of increasingly complex incidents.

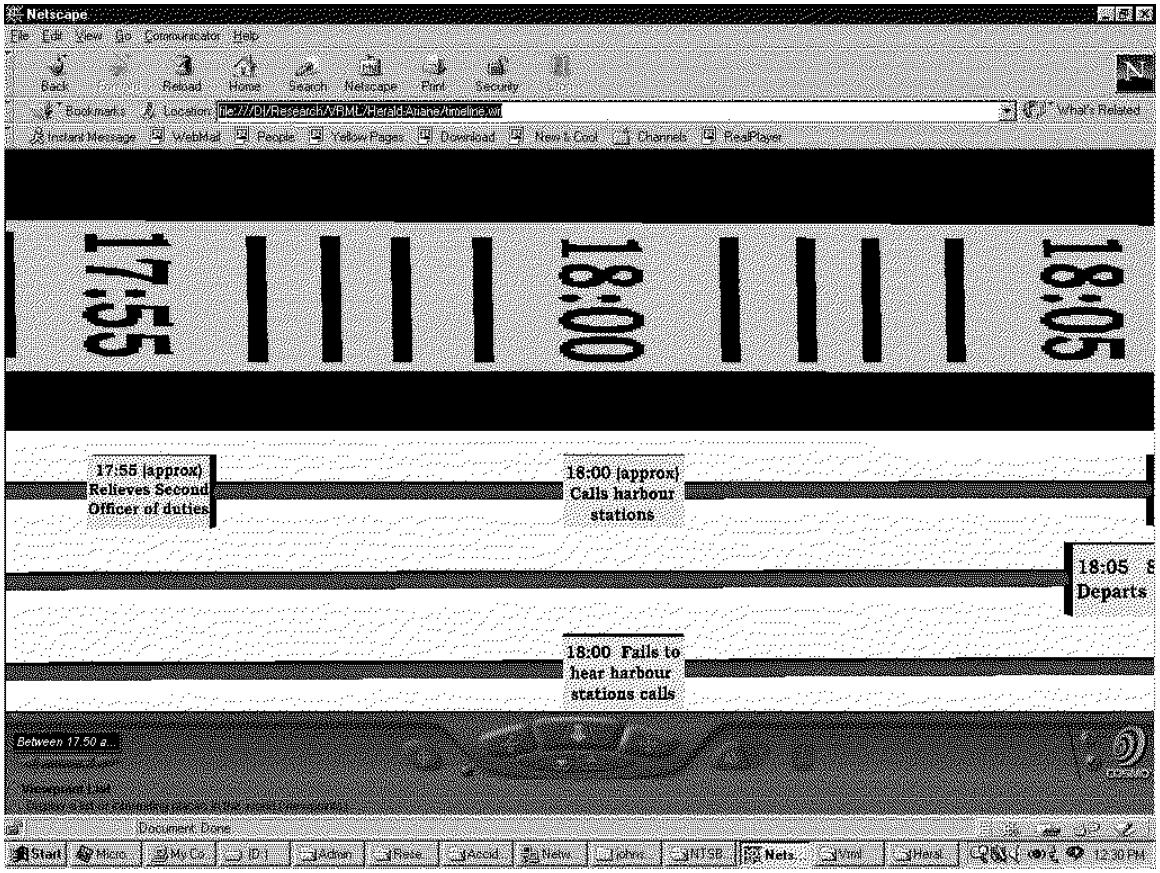
Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: file:///D|/Research/VRML/Herald-Ariane/timeline.wrl
What's Related
Instant Message WebMail People Yellow Pages Download New & Cool Channels RealPlayer
17:55
18:00
18:05
17:55 (approx) Relieves Second Officer of duties
18:00 (approx) Calls harbour stations
18:05 S Departs
18:00 Fails to hear harbour stations calls
Between 17.50 a...
Viewpoint List
Document: Done
Start Micro... My Co... D:\ Admin Rese... Accid... News... johns... NTSB Nets... Vrml Hera... 12:30 PM

Figure 8.9: Detail of Perspective Wall Using DesktopVR

8.2.3 Subjunctive Simulations

The reconstruction techniques that we have considered up to this point support declarative models of the state of an incident or linear sequences of events leading to an adverse occurrence. In contrast, this section focuses on the use of subjunctive simulation techniques for incident reconstruction. The term 'subjunctive' is used to denote the fact that many simulations are used to explore hypothetical scenarios. They provide a means of analysing what might have happened during an incident. Analysts can adjust the parameters of the model to assess the potential consequences of those changes upon the final outcome of an adverse occurrence. As we shall see, however, the reliability of any conclusions is dependent upon the degree to which the simulation faithfully recreates particular aspects of an incident. It is, typically, impossible to provide completely accurate simulations even using the physical reconstruction techniques mentioned above. In consequence, investigators must analyse which aspects of an incident must be reconstructed if we are to trust the findings of a subjunctive simulation.

Computer-Aided Engineering and Process Based Simulations

Computer-Aided Engineering tools provide one of the most powerful means of deriving interactive reconstructions of complex incidents. This software enables analysts to construct models that are intended to reflect the physical properties of the system itself. For instance, Cole and Cebon have developed a two degree of freedom mathematical model that simulates dynamic tyre forces for tractor-trailor combinations of nine articulated vehicles [172]. The results from this model have been validated in a number of empirical studies and have been used to identify combinations that yield particularly ‘strong dynamic interaction’! Others models predict the behaviour of gases within different environments. Numerical techniques derived from computational fluid dynamics have been applied to explain particular combustion processes [550]. An important benefit of this approach is that models, which are used in the design of an application, can also be used to analyse potential incidents. For example, models that explain the properties of fire resistant surfaces, for instance in furnaces, can also be used to explain why other surface fail in similar situations [769]. The dual predictive and analytic nature of such engineering models offer huge benefits, not least in the costs that would otherwise be associated with model development for incident reconstruction.

A recurring theme in this book is that there will be more than one set of events that can lead to the same adverse outcome in complex, technological systems. For this it follows that subjunctive simulations have an important role to play in searching for these alternative paths. Chapter 9 will briefly introduce techniques from abstract model checking that have been deliberately developed with this in mind. For now, however, it is sufficient to stress that subjunctive simulations can also build upon, and contribute to, training activities. For example, NASA have developed a simple interactive model of the DF88 directon finder that is used to locate aircraft Emergency Locator Transmitters [566]. Users can interact with this model to explore and practice different location scenarios. The same underlying software can also be used to during lab-based studies that are intended to explore the use of such applications in the aftermath of ‘real’ incidents. This illustrates the important point that subjunctive reconstructions, in common with the other reconstruction tools in this chapter, should also be capable of modelling the events that occur after an incident has taken place. As we shall see, more lives may be threatened by an inadequate response to an incident than are lost in the immediate aftermath of an adverse event.

A number of limitations affect the use of engineering models to explore alternative scenarios for safety-critical incidents. Many engineering models gain their predictive power by focusing on specific aspects of more complex interactions. Many of the factors that must be considered during the investigation of ‘real world’ incidents must, therefore, be excluded. For example, computation fluid dynamics can be used to model the combustion of particular materials. These models cannot, however, be applied to predict the impact that the progress of a fire will have upon key electrical subsystems during an incident. Fortunately, a number of companies have developed integrated toolsets that can be used to address such concerns. These tools enable designers to intergrate models from different engineering disciplines. The output of a conflagration model might, therefore, be used to drive predictions about the integrity of electrical subsystems. Figure 8.10 illustrates the interface to Boeing’s Easy5 tool. This software can be used to simulate systems containing hydraulic, pneumatic, mechanical, thermal electrical and digital sub-systems. Simulations can be constructed in a number of ways. For instance, functional blocks can be used together with pre-defined components that model physical elements, such as pumps, gears, engines, etc.

Other limitations stem from the implicit assumptions that must be made in orer for these models to be tractable. This is illustrated by the caveats that precede a set of results from simulator studies into aircraft icing incidents:

> “It is important to take in consideration the following assumptions and limitations of this analysis:
>
> 1. The flight conditions just prior to the upset was considered a steady state condition, meaning that all angular rates were considered small and the dynamic aerodynamic derivatives could be considered negligible.
> 2. The Power Effects (Specially the propeller slipstream effect) in the EMB-120 is very strong and for this preliminary analysis was not fully considered when calculating some aerodynamic coefficients.
> 3. The ice effects on the aerodynamic coefficients were taken from wind tunnel test results and only some Reynolds Number corrections were applied.
> 4. The flight simulation (6 DOF) is valid only up to the pusher firing angle of attack (approx. 12.5 deg). Above this angle the aerodynamic data and the effects of any asymmetric flow separation are not valid or not considered.

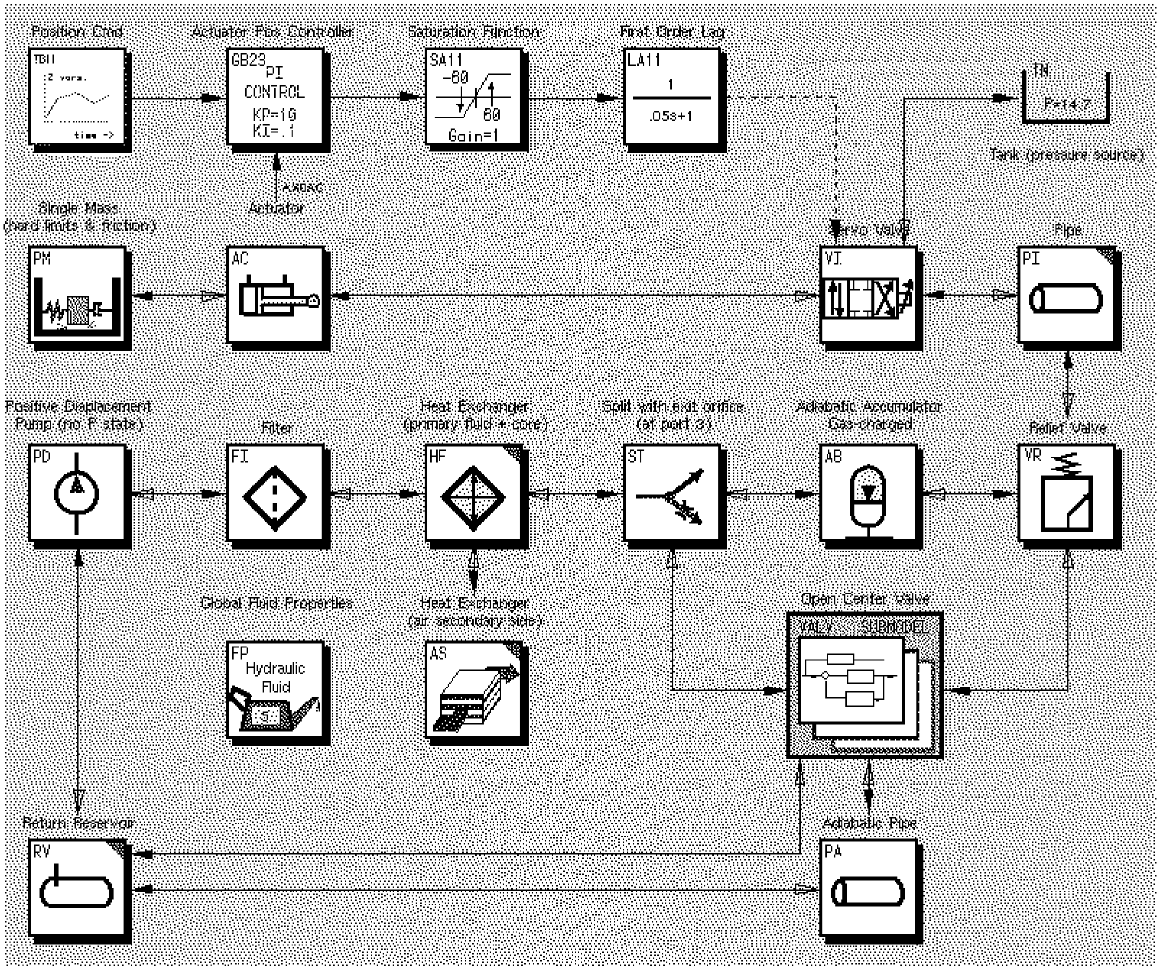

Figure 8.10: Graphical Modelling Using Boeing's EASY5 Tool

> 5. For this first preliminary flight simulation, only some aerodynamic parameters were modified and for this reason some special assumptions were made due to lack of time.
>
> All assumptions, however, were considered not relevant to this preliminary analysis." (Appendix E, [600]).

There are further limitations with existing systems. Most approaches work best for physical systems that have a linear behaviour. Systems that exhibit nonlinear and time-varying characteristics pose a considerable challenge. Multi-model systems offer a potential solution. These exploit a 'divide and conquer' strategy that represents the behaviour of a complex process in terms of a number of simpler component models. However, there can be considerable instability as a model switches between these component simulations [554]. Engineering models of human operators can be seen as a pathological example of a non-linear system. Subjunctive reconstruction systems have been developed to simulate the potential behaviour of individual operators [440]. However, these tend to be based upon high-level cognitive models rather than the more fine grained control-level simulations of engineering reconstructions. There have been a number of recent attempts to develop hybrid techniques that capture elements both of continuous control in the context of more discrete human decision making, for example in the domain of air traffic control [774]. It remains to be seen whether or not these techniques can be extended from a limited number of case studies to support the reconstruction of more complex incidents and accidents.

In contrast, discrete models of individual and group cognition have been more widely applied to simulate the events leading to safety-critical incidents. They have even been developed to model the behaviour of crowds and entire populations during emergency situations [231]. Much of this work has been inspired by the success of epidemiological modelling that often depends upon assumptions about human behaviour. For instance, Moss has recently extended this epidemiological work to derive simulations of middle management behaviour during critical incident management [551]. This work is particularly significant given that most work on simulation in incident analysis tends to focus on operational failures rather than management or regulatory behaviour. A range of similar models have, however, been developed within the field of management studies to enable users to speculate on the influences that affect key decisions [297]. The theoretical underpinning for many of these multi-agent, predator-prey models is provided by game theory. Individual operators are assumed to pursue independent goals with limited resources under conditions of uncertainty [714]. It is perhaps surprising that few of these models have been used to support incident investigation, especially given the emphasis that Reason and others have placed upon the organisational precursors to failure [701]. On the other hand, it remains to be seen whether such models actually help to analyse different patterns of individual and group performance under high-stress situations. In particular, there continue to be practical and ethical difficulties associated with the validation of these models.

Model Driven Virtual Environments

As mentioned, investigators can use subjunctive simulations to reconstruct alternative hypotheses about the course of an incident. Models of underlying physical processes can be used to replicate the 'real-world' behaviour of complex applications. Evidence that has been gathered during primary and secondary investigations can then be used to parameterise these models. If critical values are unknown then investigators can iteratively inject potential estimates for those numbers. They can then inspect the behaviour of the simulation to validate their estimates. If the modelled behaviour faithfully reconstructs key observations about the course of the incident then the values may be retained. If there are significant differences between the simulation and the observed behaviour then either the values must be revised and the test run again or questions must be asked about the veracity of the observed behaviour.

Bolte, Jackson, Roberts and McComb provide an example of subjunctive reconstruction when they describe the NTSB's use of simulations in road traffic accidents [87]. Data from event recorders can be used as input to a growing range of traffic simulation programs that model the performance of cars, buses, trucks and trains. data These devices are currently only fitted to a minority of commercial vehicles and so the software also relies upon information from witness statements to deduce probable driver inputs. Physical evidence, including the final resting point of the vehicle and any resultant damage, can also be introduced as parameters to the current generation of reconstruction software. All of this information helps to produce crash pulses. These graphs describe the likely acceleration profiles that would be required in order for each vehicle to end up in their final positions with the degree of damage that was recorded. Figure 8.11 illustrates the crash pulse that the NTSB can obtain from road-traffic simulation software.

A collision between a van and a train at Wagner, Oklahoma can be used to illustrate the role that such graphs play during an incident investigation [87]. Witnesses in the van reported that the driver stopped before proceeding over the railway crossing. However, the train engineer contradicted this statement. The train had an event recorder on board and data was obtained from this to determine that it was travelling at approximately 46 miles per hour when the collision occurred. The final position of the van was surveyed in the aftermath of the incident. Data was also collected about the damage that the train had inflicted in the collision. "Reconstructionists" at the NTSB then varied the speed of the van at the point of impact to determine the most accurate trajectory and the related initial speed of the van. This illustrates the way in which subjunctive simulation techniques can be used to explore alternative hypotheses about the events leading to an incident; the van did or did not stop before the collision.

The graphical format used to represent crash pulses, illustrated in Figure 8.11, has a number

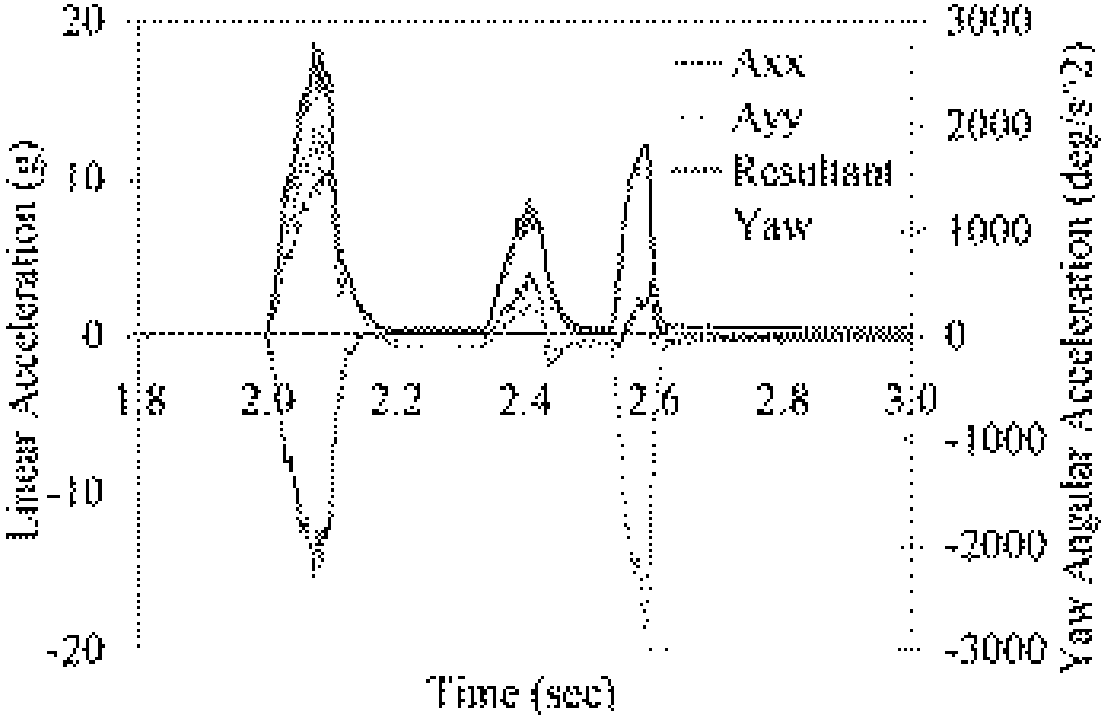

Figure 8.11: NTSB Simulated Crash Pulse Of School Bus and Truck Colliding

of benefits. In particular, it provides a precise numerical profile for vehicle acceleration during an incident. However, it can be difficult to use such representations to communicate the key findings from an incident investigation. This is a particular problem when multi-disciplinary investigation teams must agree about the likely course of events. In consequence, most simulation software provides a range of alternate visualisations. Investigators can use these to show their colleagues what particular calculations imply about the likely course of events leading to an incident. This is illustrated by the model-based virtual reconstruction shown in Figure 8.12. The rendering software that is used to generate this image is very similar to that used in Figure 8.5. However, there are a number of important differences between simple animated simulations and the subjunctive systems illustrated in this section. The former systems rely upon ad hoc scripting techniques to develop a single model of an incident. As mentioned previously, this is often used to derive a linear sequence of images that are edited to provide a movie that can be played during an incident reconstruction. In contrast, subjunctive techniques rely upon physical models of application processes. As a result, it is possible to explore multiple alternative hypotheses about the course of events in an incident by altering the parameters of those physical models.

Simulations, such as Figure 8.11, are not the end-point for inicident reconstruction. Once an accurate simulation has been developed to model vehicle performance, it is then possible to reconstruct occupant kinematics. This is important, especially in road traffic accidents, because it can be used to assess the degree of occupant protection that the vehicle afforded. It is important not to underestimate the importance of such feedback for the design of future safety features.

Operator behaviour must be explicitly considered as a parameter to most subjunctive simulations. This raises many problems. In systems without data recorders, investigators must typically rely upon eye-witness statements. As we have seen with the Wagner incident, these can be contradictory, biased and partial. It is difficult to envisage a set of circumstances in which both the train driver and the van's occupants were correct in their accounts. However, many incidents occur during periods of

Figure 8.12: NTSB Simulation of Motor Vehicle Accident, Wagner Oklahoma

operator inattention or fatigue. In consequence, they may have genuine difficulty in recalling their actions. Most traffic accidents last less than 0.10 seconds. They happen at a speed that makes it difficult for witnesses to accurately comprehend 'vehicle interactions' [87].

The problems of modelling operator interaction in subjunctive simulations are considerably eased if information can be obtained from automated data logging systems. For example, Figures 8.13 and 8.14 illustrate how this data can be used to drive biomechanical simulations of the operator's actions [609, 605]. These simulations represent the rudder pedal positions and leg orientations of the crew. Although this cannot be seen from the still image, the colour of the manikin's leg indicates the force output. For example, blue as used to indicate that no force is being applied to the rudder pedal. Yellow indicates a normal force application while red indicates a larger force than would normally be needed during a flight. Given individual human variability and the sparse data that is often obtained from logging systems in the aftermath of an incident or accident, these models must often be treated with caution. This is evident in the caveats that accompany the NTSB's use of these simulations:

> "These simulations were developed as an educational aid although, whenever possible, the scaling and motions of the manikin and cockpit control were modeled after those of the Boeing 737 event being studied." [605]

Figure 8.14 provides an overview of crew interaction in the cockpit. The NTSB have used annotated similar reconstructions with to include excerpts from cockpit voice recorders as subtitles to the model based simulations. This enables investigators to follow key communications as they watch reconstructs of the crews' interaction with their controls. Such techniques are interesting for many

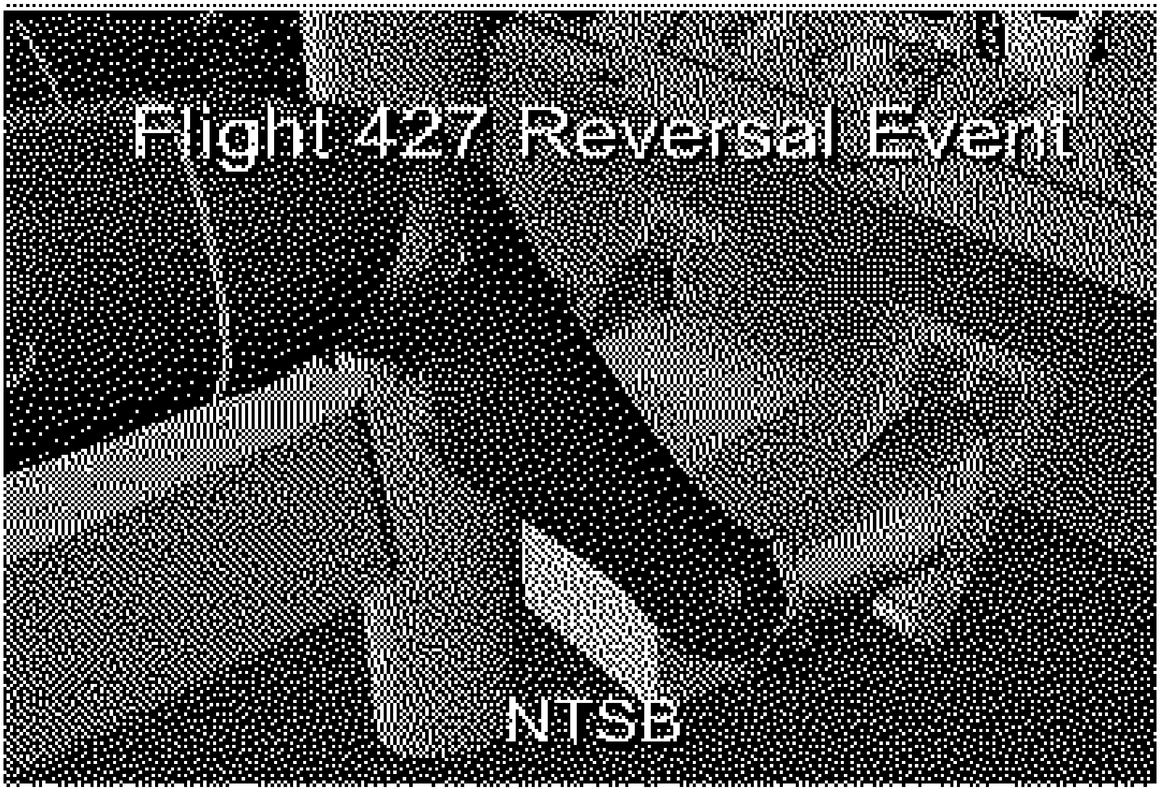

Figure 8.13: Biomechanical Models in NTSB Incident Simulations (1)

reasons. In particular, the public release of these simulations has not created the same controversy as fears about the public release of footage from cockpit video recorders [605]. The US Air Force recognises the distinctions between these different media:

> Animations made from recorder data are not privileged as long as they do not contain (the investigation board's) analysis or input. If the actual audio voices of the mishap crew are incorporated into the animation, simulation or re-enactment videotape, the tape is not releasable due to the privacy interests of the crewmembers or their surviving families." [794]

There are numerous limitations with the subjunctive techniques that are described in this section. In particular, it can be difficult to gain the data that is necessary to derive the models of application components. This means that the resulting visualisations, such as that shown in Figure 8.12, may ultimately rest on little more than guess work. For instance, road traffic simulations typically require information about both drivers' behaviour. This includes steering angles, brake and throttle settings, gear selection, use of engine braking. Some models also expect information about the use of lights, of direction indicators cruise control settings and even wipers or warning horn. In addition, the models require information about the pre-impact speed, engine revs. per minute, acceleration history, braking efficiency activation of anti-lock braking mechanisms etc. If these parameters are not provided then systems either fail to produce a simulation or they exploit default assumptions that may not reflect the conditions that held during a particular incident. Such problems emphasise the point that many subjunctive simulations reconstruct a probable or possible course of events. They do not provide a definitive and unambiguous account of most safety-critical incidents.

Figure 8.14: Biomechanical Models in NTSB Incident Simulations (2)

Probabilistic Simulations

The use of subjunctive simulation techniques is hampered by the problems of data acquisition. Incident recorders are often unavailable. Even if this source of data is available, they often only provide a small subset of the parameters that are required by many reconstructions. The situation is not as problematic as it might first appear. In particular, performance metrics can often be recruited to support data recorders and eye-witness information. Even when these metrics are not available, standard reliability techniques can be used to estimate the potential failure rates for particular pieces of equipment. For example, Table 8.6 presents availability data for the Display Channel Complex computers that formed an important component of the US Air Traffic Management infrastructure. This is calculated as follows:

$$Availability = \frac{Mean\ Time\ Between\ Failure}{Mean\ Time\ Between\ Failure\ +\ Mean\ Time\ to\ Repair}$$

Although Table 8.6 focuses on availability, the same techniques can be used to model more general process behaviour including network loading and resource scheduling. The key point is that the resulting stochastic models mirror the probabilistic behaviour of system components. They can, therefore, be used to determine the behaviour of any simulation in situations where data recorders were unavailable or where their readings provide suspect results.

Monte Carlo techniques provide one relatively simple means of using availability data, such as that shown in Table 8.6 to support subjunctive simulations. This proceeds by generating a random number in the interval between 0 and 1. If the number is less than the availability rate for that component then that component is assumed to be available in the next interval. However, if the number is greater than the availability rate then the component is assumed to fail. For instance, the availability of the Display Channel Complex at the Washington centre in the first quarter of 1994 was 0.9086. If the random number generator produced 0.5 then the simulated Display Channel Complex

Quarter	Chicago	Cleveland	Forth Worth	New York	Washington
Ql, 1992	0.9877	0.8265	0.8899	0.9349	0.9403
Q2, 1992	0.9720	0.9716	0.8059	0.1143	0.9845
Q3, 1992	0.9413	0.9384	0.6984	0.9563	0.9768
Q4, 1992	0.9591	0.9356	0.6566	0.3769	0.9409
Ql, 1993	0.7373	0.9324	0.6628	0.8258	0.8794
Q2, 1993	0.9361	0.9744	0.9294	0.0000	0.9548
Q3, 1993	0.8143	0.9537	0.9011	0.7636	0.9708
Q4, 1993	0.6185	0.9646	0.7397	0.7846	0.9552
Ql, 1994	0.4942	0.9739	0.9458	0.6636	0.9086
Q2, 1994	0.7521	0.8741	0.8257	0.7916	0.8480
Q3, 1994	0.9608	0.9504	0.6955	0.9177	0.9190
Q4, 1994	0.9526	0.9670	0.7996	0.7738	0.9520

Table 8.6: Availability of US ATC Display Channel Complex Computers [590]

would not fail during that quarter. However, if it produced 0.9100 then the simulated component would fail. If the availability increases then it becomes less likely that the random number will be greater than this revised figure and so the component becomes less likely to fail in any simulation. If the availability decreases then it becomes more likely that the random number will be greater than this figure and so the component becomes correspondingly more likely to fail in any simulation.

The previous paragraphs have presented a simplification of the probabilistic techniques that investigators can use to support subjunctive simulations [27]. The key point, however, is that these models help to ensure that reconstructions are based upon observed behaviours even when there may not be direct data about the course of an incident. In particular, information about previous failures can be used to bias simulations towards particular traces of events. For example, if a particular installation had continual problems with their Display Channel Complex then the availability figures would be considerably reduced during reconstructive simulations.

There may seem to be a paradox in the previous discussion. We have argued that subjunctive simulations help investigators to explore alternative hypotheses about the course of an incident. However, this section has argued that reliability data and probabilistic simulations can be used to ensure that the behaviour of a simulation is narrowly based upon the observed behaviour of particular components. This apparent paradox can be resolved by emphasising that these techniques still enable investigators to identify a range of 'what if' scenarios. For example, Monte Carlo techniques inevitably support some of this exploration because different random numbers should be generated on each pass. It should be apparent, however, that rare events will still be very difficult to simulate under this approach. In simple terms, if a component has an associated availability rate of 0.99 for a particular time period then there is 1 in 100 chance that it will fail in any particular run of a simulation. If an investigator wished to focus on those scenarios in which this component were known to fail then they might make a corresponding reduction to the associated availability of that component. The danger here is that such changes might not accurately reflect the availability of that component during an incident. The more changes that are made, the further an investigation moves from any available reliability data. In consequence, it is important that investigators keep a log of the sources of data that are used to drive any simulation together with a detailed justification of any changes that are made to such data.

Probabilistic simulations support subjunctive reconstruction in a number of further ways. One of these approaches can be illustrated by changes in the maintenance procedures for the Display Channel Complex equipment, mentioned above. Approximately five years before the figures in Table 8.6 were published, one of the US ATM sites conducted a hazard analysis for these components. This raises a number of concerns. For example, the standard maintenance practice was to immediately intervene whenever one of the three computing elements failed. The review demonstrated that this might disable the two remaining computing elements and thereby jeopardise service provision. In consequence, a revised plan required that the engineers wait for a low traffic period to begin repairs. In consequence, there was a deliberate decision to continue operating with less-than-full redundancy. Local technicians and managers believed that this is 'less risky' than beginning immediate repairs. Table 8.6 presents availability data for a fully redundant Display Channel Complex. During subjunctive simulation, it would be relatively straightforward for investigators to replace this with information with availability data from sites that operated the revised maintenance policy. Such techniques illustrate one of the ways in which simulations can be used to generalise beyond the specific circumstances of a particular incident to examine other failure scenarios.

These applications of probabilistic simulations must be considered against a number of limitations. Chapter 3 has already described the practical and theoretical problems that are associated with the stochastic modelling of software failures. These concerns form part of a wider scepticism about the probabilistic modelling of safety-critical incidents. For instance, Chapter 2 has described Wright's recent work in the railway industry [874]. She has shown that previous information about a large number of relatively low criticality failures often provides few hints about the likely state of a system during higher criticality incidents. These caveats have recently persuaded a group of researchers under John Fox at the Imperial Cancer Research Fund to look at *possibilistic* simulations. The intention is not simply to explore what is likely but rather to determine what is feasible given a particular model of the system. The obvious drawback to this approach is that many possible scenarios may lead to a particular incident. Possibilistic simulations, therefore, depend upon arguments about the plausibility of a particular scenario given what is known about a set of events. The emphasis is, therefore, placed upon a more general argument of plausibility rather than a search to quantify reliability estimates. These ideas need to be more fully developed. In particular, more work needs to determine the appropriate form for convincing and plausible arguments about the veracity of a simulation. As we have seen, many reconstruction techniques offer ample opportunities for biasing simulations towards particular viewpoints about the causes of complex failures. There are, however, strong links between this approach and, for instance, the increasing use of model checking in safety and reliability analysis. This approach models an application in terms of a number of possible state transitions. These transitions describe how the state of a system can change over time. Analysts can then specify properties or theorems that they would like to hold over the system. Automated model checking tools will then explore the possible states of the system to determine whether or not the property actually does hold for this model of the system [192].

A number of further issues complicate the application of subjunctive simulation techniques for incident reconstruction. Arguably the most important of these surrounds the modelling of team-based interaction during adverse occurrences. Previous paragraphs have explained how cockpit voice recorders and flight data recorders can be used to integrate human factors observations into physical models of safety-critical incidents. However, we have not explained how reconstructions might be developed when such data is unavailable. Probabilistic techniques, such as Monte Carlo simulations, can be used to address this limitation. System operators can be asked to interact with a simulation as it reproduces a probable failure scenario. Operator behaviour can be monitored over successive runs to identify a range of possible responses. This is, however, a partial solution. Operators often alter their behaviour if they know that they are being monitored in the aftermath of an incident. Such problems are exacerbated if operators must interact with computer-based simulations rather than with their 'real world' counterparts. Moving a desktopVR model of a train around a simulated track is hardly comparable to the physical and mental processes involved in driving a real locomotive. Similarly, it can be extremely difficult to reconstruct the working pressures and team-based interactions that characterise most working environments.

Single and Multi-User Simulations

It can be difficult for individuals to recall their actions immediately before an incident. Even if they can remember, personal, organisational and social pressures can prevent operators from rendering accurate accounts of their behaviour. In consequence, it is important that investigators have some means of reconstructing alternative hypotheses about the role of human intervention in an incident.

There are two main approaches to this form of simulation. The first is to study human interaction with computer based reconstructions, in the manner described above. The second is to derive computer-based simulations that attempt to model human intervention with complex systems.

The following list summarises some of the problems that arise when studying operator interaction with simulated systems. These problems stem partly from the difficulty of staging such reconstructions and partly from interpreting the observations that can be derived from them:

- *problems in obtaining access to appropriate staff.* In the aftermath of an incident, it can often be difficult for investigators to obtain the cooperation of operators who are willing to participate in simulation studies. Such volunteers must possess the necessary skills and experience to take part in the studies. They must agree to have their performance monitored and recorded while they interact with a simulation, even over prolonged periods of time. They must be prepared to 'suspend disbelief' when mock-ups replace application processes. Even if such individuals can be found, it is often important to secure the support of trades unions and other forms of worker representation before many workers will participate in such studies. This usually involves assurances about the ultimate use of any data that is to be derived from the studies;

- *difficulty of simulating individual factors.* The list of requirements in the previous paragraph can have a paradoxical effect on participant selection for simulator studies. They imply that potential 'subjects' must have an active interest and involvement in safety issues. Individuals who are prepared to have their interaction monitored over prolonged periods of time may not provide results that can be generalised across the rest of the workforce. Even if such biases can be counteracted, it seems unlikely that investigators will be able to address the diverse range of behaviours that characterise individual responses to the stress and uncertainty of many incidents. In other words, simulator studies may only provide a partial glimpse of the operator behaviours that might have occurred during a particular failure;

- *difficulties in repeat simulations.* Operator interaction can be observed during several trials with the same scenario. This helps to reduce any nervousness during an initial observation. It can also help to reduce any effects that stem from differences between a simulation and the 'real world' system. However, operators may gradually transfer knowledge gained from previous trials to guide their interaction with subsequent simulations. This not only applies to repeated trials with the same scenario, knowledge can also be transferred between different simulations. It can, therefore, be argued that operator interaction can provide less and less realistic results as the number of trials increases. Eventually there is a danger that boredom and fatigue may provoke behaviours that would not be apparent in other circumstances;

- *difficulty of simulating contextual events.* Chapter 3 briefly introduced the range of performance shaping factors that can affect operator behaviour. These can include heat, noise, light-levels, individual fatigue, alcohol and drugs etc. In anonymous reporting systems it is often impossible to determine whether ot not these played a significant role if they are not explicitly mentioned in an incident report. Even if investigators can interview operators, they may not be aware of the extent to which these factors have influenced their actions. Assuming that investigators have identified the importance of heat, noise etc, they then face the task of accurately recreating these influences in a manner that is ethically acceptible. This is less easy than it might at first appear. In some studies, operators exhibit a relatively small number of errors under ideal conditions. In consequence, investigators have started to bombard them with noises, flashing lights and other forms of distraction so that the simulations become parodies of the environments that they represent;

- *difficulty of provoking rare behaviours.* Many of the operator behaviours that exacerbate or cause incidents are extremely rare. In most cases, systems are designed to mitigate the consequences of such actions. Individuals may also be preselected and trained to reduce the likelihood of such intervention. In the aftermath of many incidents, operators can also be sensitised by rumours of their colleagues' intervention. This may make individuals even less likely to repreat those previous failures. All of these factors may force investigators to run many hundreds of trails before evoking a response that might have contributed to an adverse occurrence;

- *difficulty of interpreting observations of simulated behaviour.* Even if it is possible to overcome many of the practical barriers that frustrate the observation of user involvement with incident simulations, there are still many problems associated with interpreting the behaviours that are identified. For example, the same user may interact in one way during one trial but then interact in a quite different manner during a subsequent run. Alternatively, two different operators may react in quite different ways to the same simulation. A range of techniques can be used to examine these differences. For instance, individuals can be asked to explain the motivation for their decisions as they interact with a simulation. However, this process of introspection may force them to re-examine their decisions in a way that would not occur during normal interaction with the system. Alternatively, cognitive modelling techniques can be used to represent some of the psychological processes that might motivate individual interaction [122]. However, it can be difficult to ensure any form of inter-analyst consistency using this approach. There are often considerable disagreements about the underlying mechanisms that provoke particular operator actions during simulator studies;

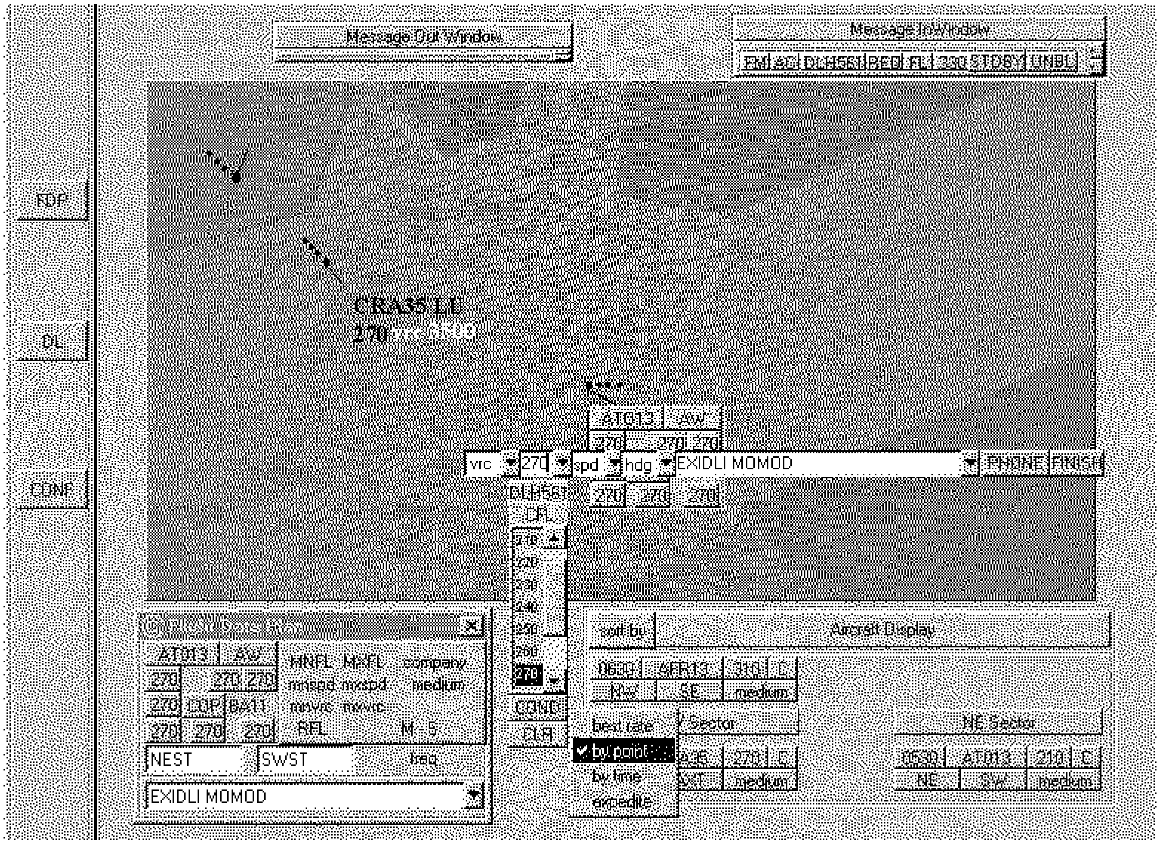

Figure 8.15: Multi-User Air Traffic Control (Datalink) Simulation

- *difficulty of simulating team dynamics.* Previous paragraphs have mentioned the problems of simulating group based interaction. Each of the previous items in the list could be re-written specifically in terms of the difficulties associated with team based simulation. For instance, it is difficult to underestimate how complex it can be to interpret those factors that motivate particular group-based activities. In many incidents this involves a recursive dependency in which my actions can be influenced by which I think that you are thinking about my

> current intentions [405]. As mentioned in Chapter 3 team-based interaction can be dominated by particular individuals. This may or may not reflect what actually happened during a particular incident. It can often be difficult to recreate the balance of skills and personalities that contributed to an incident. Similarly, the problems of securing operator participation are exacerbated by the pragmatic difficulties of ensuring that different operators are all available at the same time. There can also be significant costs associated with the development of multi-user simulations. Figure 8.15 illustrates an interface that has been produced using one of a number of environments that are intended to reduce these costs [719, 720]. This multi-user interface builder has the added benefit that it can be used in conjunction with the model checking tools that were mentioned in the previous section. This example shows an en-route controller's view of a datalink air traffic control system. They can communicate with other controllers and with individuals playing the roles of the aircrews in their sector. The expense involved in running a simulation session with qualified operators should be obvious. The complexity in implementing such a system and then linking it to accurate traffic models should also be apparent.

Given all of these problems, a growing number of researchers have turned to alternative means of simulating operator interaction with safety-critical systems [727]. Rather than creating environments that are intended to enable human operators to replicate the events leading to failure, this approach extends the scope of computer-based simulations. In particular, research teams have developed a number of computer-based simulations that are intended to recreate individual behaviour. Some of these models have also been 'hooked up' to other simulations of system behaviour in order to observe potential interaction. Again these techniques are still in their infancy. For instance, NASA Ames research centre recently reported on the APEX project that was specifically intended to make Human-Machine system simulation a practical engineering technique' [703]. This project has identified an architecture for future systems. The interpretation of auditory and visual input helps to determine the contents of the operator's working memory. This, in turn, influences the process by which an individual selects their actions for intervention. These actions are executed through components that model gaze and attention as well as vocal commands and the operator's gestures. At present, this work focuses on modelling the behaviour of individual operators. It has the important advantage of explicitly representing theories about the factors that motivate operator behaviour. Although in the future it might be possible to compose several of these models to simulate group interaction, it remains unclear how this might be achieved in practice. I am also unaware of any attempts to apply this approach during an incident investigation.

8.2.4 Hybrid Simulations

This second half of this chapter has identified a number of different simulation techniques that investigators can use to reconstruct the events leading to safety-critical incidents. In identifying these different approaches it has, however, been necessary to introduce what are often arbitrary distinctions. For example, subjunctive simulations can be used during the initial stages of an investigation. Once analysts have reconstructed a likely course of events, this may then be used to produce animated simulations. These prevent viewers from exploring the alternative options that are available to incident investigators. In contrast to subjunctive techniques, users are simply presented with a linear sequence of images that illustrate a particular perspective on an incident. There are further complexities. For instance, declarative reconstructions can be used to model the physical environment in which an incident occurs. However, they provide little information about the way in which particular events contribute to an adverse occurrence. They must be used in conjunction either with text-based time-lines or with abstract simulations of critical events. It can, therefore, be argued that most reconstruction tools exploit hybrid combinations of the techniques that we have reviewed.

Figure 8.16 provides an example of a hybrid approach to incident reconstruction. This design was developed by Gilles Le Galo as part of a proposed simulation tool for incident investigation throughout European air traffic control [423]. As can be seen, the top portion of the display uses data logs to provide real-time reconstructions of the controllers primary information displays. Below this there is an area that reproduces the flight strips that were being used by the controller while they were interacting with this information. This is an interesting use of hybrid reconstruction

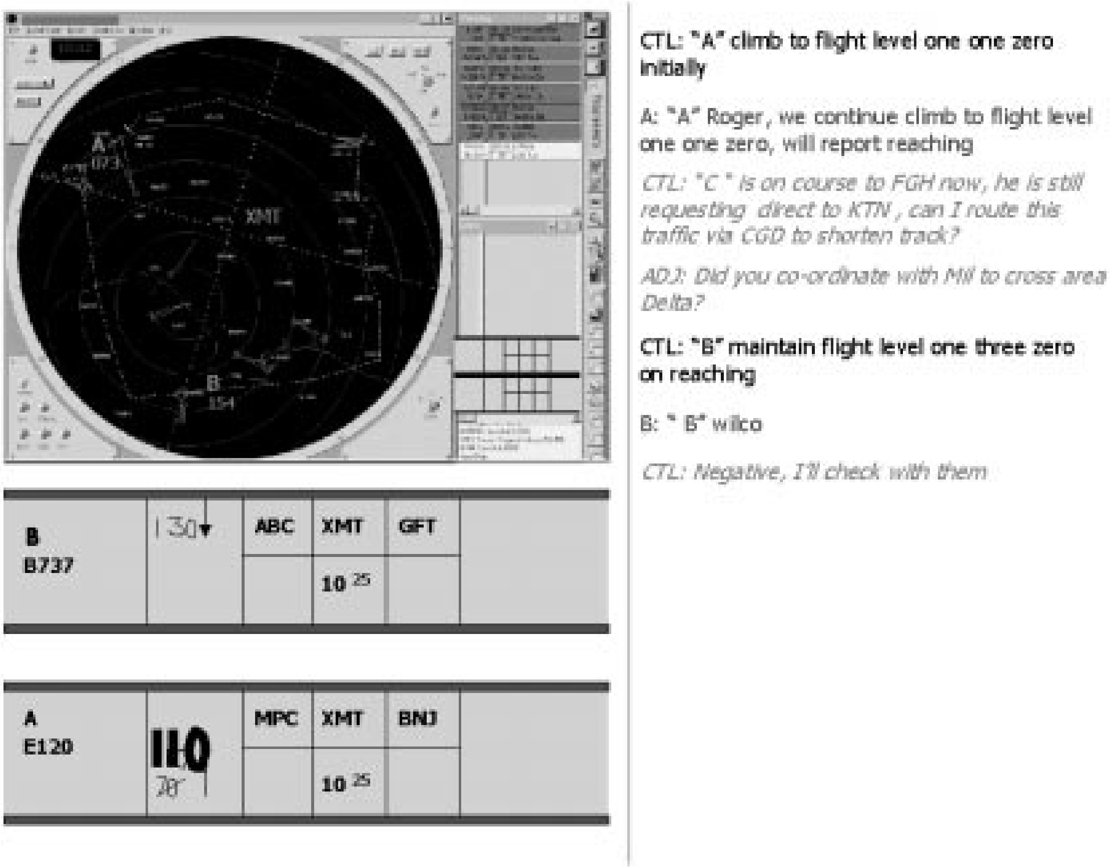

B B737	130▼	ABC	XMT	GFT
			10^{25}	

A E120	110 70r	MPC	XMT	BNJ
			10^{25}	

CTL: "A" climb to flight level one one zero initially

A: "A" Roger, we continue climb to flight level one one zero, will report reaching

CTL: "C " is on course to FGH now, he is still requesting direct to KTN , can I route this traffic via CGD to shorten track?

ADJ: Did you co-ordinate with Mil to cross area Delta?

CTL: "B" maintain flight level one three zero on reaching

B: " B" wilco

CTL: Negative, I'll check with them

Figure 8.16: EUROCONTROL Proposals for ATM Incident Simulation

techniques because it contains elements of model based simulation; through the regeneration of the primary information displays. It also supports the electronic reconstruction of the controllers' physical environment by presenting scanned images of the relevant flight strips. Finally, the transcript on the right is used to keep track of the radio communications between controllers and the aircrews in their sectors. The vertical, linear presentation of this information is reminiscent of the time-line annotations presented in previous sections.

[211]

There are numerous further examples of this hybrid approach. Various simulation techniques can be supported within a single system, as shown in Figure 8.16. It is, however, more common to use a range of different simulation tools to satisfy different modelling requirements during an incident investigation. This is shown by the range of approaches that are illustrated in Figure 8.17. The image shows a number of different approaches that together contributed to a US National Crash Analysis Center study of ankle injuries in automobile incidents [211]. The image on the top left of Figure 8.17 is derived from a finite element model of joint behaviour. This mathematical technique can be used to account for aspects of joint behavior, muscle tensioning, and injury potential in a high speed impact. The second image on the top of this figure illustrates the process of direct physical measurements that were used to construct a model of the environment in which an injury might occur. These measurements were used to produce the wireframe model on the top-right and the rendered image on the bottom-right of Figure 8.17. Finally, the image on the bottom left shows how a physical model of the lower leg can be used to analyse the causes of injury in an incident. Such direct measurements can be used to validate computer-based simulations of an incident. Mathematical models, rendered visualisations, direct measurements and physical reconstructions all contribute to

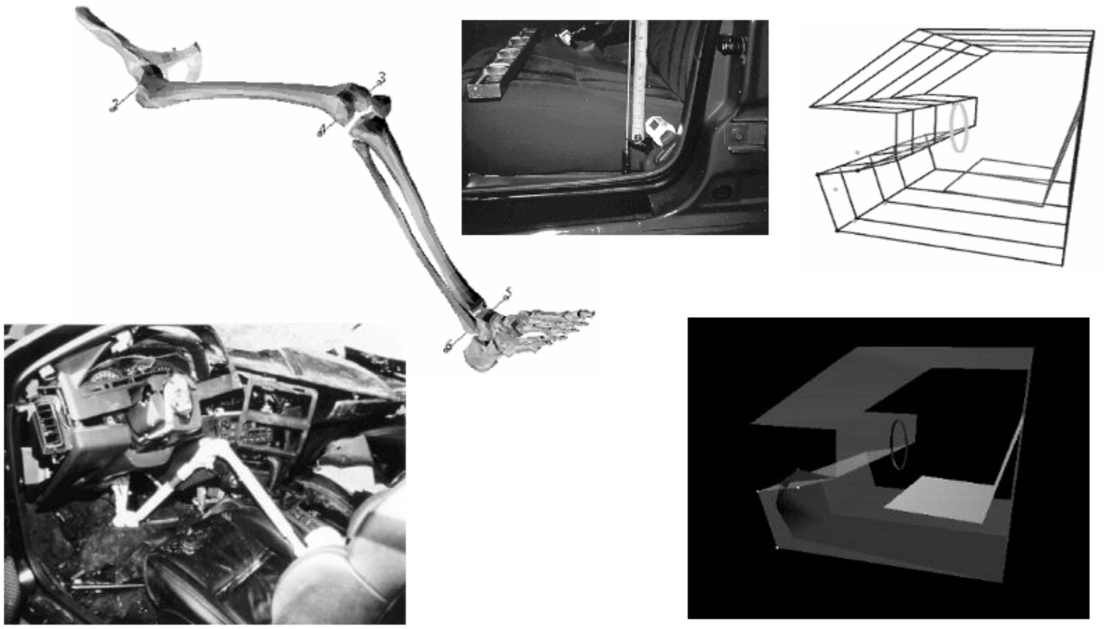

Figure 8.17: US National Crash Analysis Centre's Simulation of Ankle Injury in Automobile Accidents

increase confidence in the investigators' findings. The authors describe this integration of detailed incident investigation techniques and computer modelling as 'unconventional'. It is also, arguably, essential if we are to validate the products of computer-based simulation.

This integration of computer-based modelling and incident analysis has uncovered a number of important results. For example, the National Crash Analysis Centre's study helped to highlight the importance of muscle tension in lower leg injuries. The simulation predicted that greatest injuries occurred when the leg muscles were relaxed, as might be the case if an individual was not anticipating an incident. However, when the leg is tensed it can act like a stiffened beam so that the greatest load must be absorbed by the weakest element. In consequence, severe ankle injuries are likely if the driver is aware of a potential impact. As we have seen these results were both informed and validated by incident investigations. This could not have been done using other forms of reconstruction. "Dummies require adjustable joints and repeated recalibration while cadavers require some form of joint locking device." [211]

This chapter has focussed on computer-based simulations as a means of reconstructing safety-critical incidents. This is justified by the increasing use of these tools in many different application domains. As systems are developed to support the analysis of increasingly complex accidents, the same technology is gradually also being recruited to support the reconstruction of near-miss incidents. It is important to emphasise, however, that many of the techniques embedded within computer-based reconstruction tools are extremely general. It is possible to exploit similar simulation techniques using more conventional, pencil and paper based approaches. For instance, Figure 8.18 illustrates a hybrid approach to incident reconstruction [48]. This is based on materials that were presented by investigators from Australia's Marine Incident Investigation Unit. An engineering model of the piping system for the on-board generators can be printed together with an overview of the layout of the vessel and direct photographic evidence about the state of key components within the system.

It is important to emphasise that paper-based simulations tend to focus on declarative approaches, such as the maps and diagrams shown in Figure 8.18. It is less easy to reproduce the dynamism that characterises both subjunctive and animated reconstructions. There are, however,

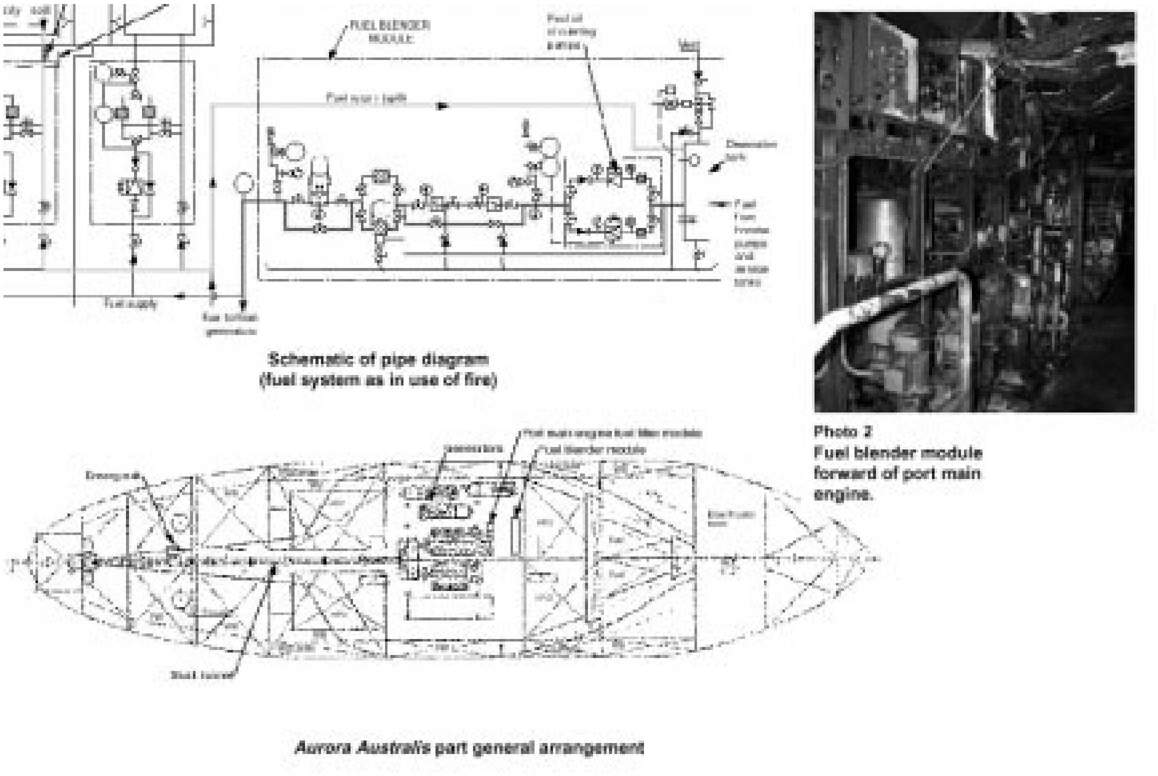
FUEL BLENDER MODULE
Fuel supply
Schematic of pipe diagram
(fuel system as in use of fire)
Photo 2
Fuel blender module
forward of port main
engine.
Aurora Australis part general arrangement

Figure 8.18: Integration of MIIU Plans, Models, Maps and Photographs

some notable exceptions to this general remark. For instance, Figure 8.19 illustrates how the NTSB have used images from desktopVR simulations to illustrate paper-based documents [596]. This image also demonstrates that many previous distinctions between paper-based and electronic reconstructions are becoming extremely blurred. Figure 8.19 shows the electronic version of a paper based incident report. It can be downloaded from the NTSB's web site (www.ntsb.gov). This electronic version of a paper based report was, in turn, generated using an electronic simulation. This translation between electronic and paper-based media is currently not well handled by many investigation agencies. It produces numerous ironies and inconsistencies. Readers who access the NTSB's website can only view the static image that is presented in the paper-based report. This misses an important opportunity to exploit the presentation techniques that are supported by computer-based, simulations.

8.3 Summary

This chapter has introduced the distinction between reconstruction and causal analysis. The former is intended to identify what happened during an incident by simulating the flow of events that led to a near-miss. The latter is intended to explain why those events occurred in the first place. The previous sections, and the next chapter, focus on reconstruction. The intention is to provide a broad overview of the tools and techniques that can be used to piece together evidence for an investigation to develop a coherent account of the course of an incident. This chapter has focuses on computer-based simulations. This is justified by the increasing use of these systems to support incident and accident investigations in many different application domains. In contrast, the next chapter looks more narrowly at the use of textual and graphical notations to model the events leading to failure. These techniques do not require computer support.

It may seem paradoxical to discuss computer-based simulations before paper-based techniques.

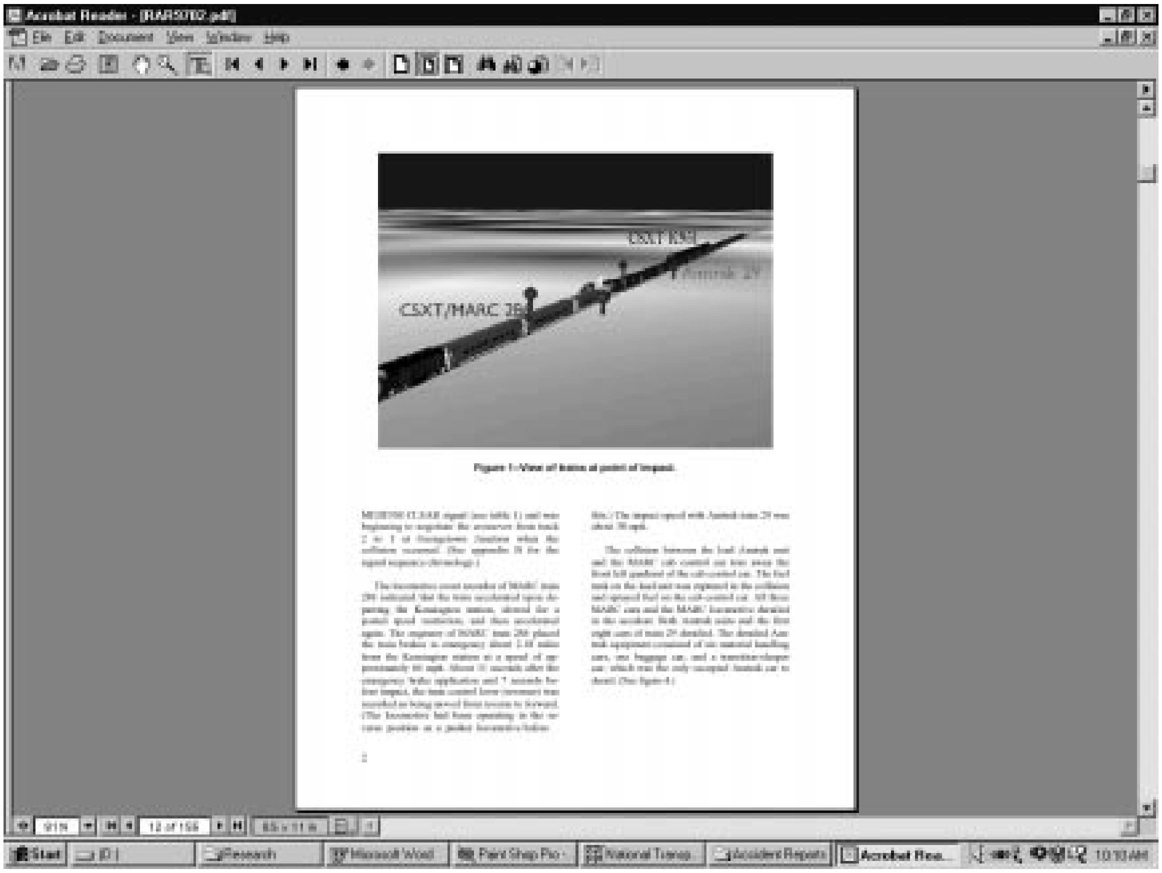

Figure 1–View of trains at point of impact.

[illegible]

[illegible]

2

Figure 8.19: NTSB Use of Simulations in Incident Reports

The intention is, however, to first show what is possible using existing software and then to explain why many investigators retain more 'conventional' techniques. Many of the systems that have been introduced in this section are still not widely used. There are many reasons for this and these can be summarised as follows:

- *The financial costs and training overheads are prohibitive.* Many local and regional reporting systems lack the resources to invest in computer-based simulation techniques. There are insufficient funds to acquire and maintain the necessary hardware and software infrastructure. In consequence, many computer-based tools are never applied beyond demonstrator projects. Those that do are most often applied within accident investigations where public and regulatory concern motivate greater investment. These objections are, however, losing their weight. Many simulation techniques, such as QuicktimeVR and the model-based VRML, have been derived from low-cost mass market applications. They can be develop and run on standard PCs and the software is designed to be used by members of the general public;

- *Individual and organisational opposition to innovation.* The reluctance to exploit computer-based simulations can perhaps be explained by the natural conservatism that characterises many investigators. Others have chosen to use more provocative terms. The Rand report has pointed out that having relatively lengthy experience requirements provides strengths and weaknesses for incident investigation [482]. On the one hand, it helps to ensure a detailed first-hand knowledge of an application domain. On the other, it can create a reluctance to exploit novel technologies. This may also apply at an organisational level. The same report criticises the NTSB's lack of any coherent plan to exploit technological innovation in accident and incident investigations. Such comments provoke strong reactions. The Rand report's comments are undoubtedly true for many investigator bodies. It seems paradoxical, however, that some of these criticisms should be applied to an organisation that has also shown a strong desire to innovate. Many of the simulations in this chapter have been inspired by the NTSB's work in this area;

- *Simulations forcer investigators to consider a small subset of incident scenarios.* There are a range of more theoretical objections to the introduction of computer-based simulations. For instance, the visual impact of many computer-based simulations can have an unwarranted effect upon some investigators. This is particularly important given the difficulty of validating the backwards reasoning that drives many simulations. Investigators are often required to estimate the forces or other properties necessary to incur the consequences that are observed in the aftermath of an incident. Unfortunately, the same observed set of outcomes can be derived from many different precursors. This creates particular concerns if simulation tools focus investigators' upon a small number of these potential scenarios.

- *Simulations can encourage encysting and may limit the scope of an investigation.* By 'encysting' we mean that investigators may spend so much of their available resources on developing an elegant simulation that they neglect other aspects of their analysis. This effect is compounded by the fact that some aspects of an incident can be extremely difficult to reconstruct using current tools. For example, it can be hard to reconstruct some meteorological conditions using many existing systems. Similarly, it is hard to envisage ways in which current tools can be used to animate the events leading to managerial and regulatory failures;

- *Reconstructions can raise as many questions as they answer.* Ultimately, simulation tools are only useful if they support the wider objectives of the people and organisations who use them. There is, arguably, a perception that investigators are adequately supported by the pencil and paper-based techniques that are discussed in the next chapter. If Perrow [675] and Sagan [718] are correct then this situation will change as incidents and accidents reflect the increasing complexity of modern technology. It does not, however, follow that computer-based simulations provide an appropriate alternative to these conventional techniques. For instance, current text-based time-lines enable investigators to work at a level of abstraction that is appropriate to the stage of their investigation. Crude sketches, which are produced during an initial enquiry, can be developed and extended as more information becomes available. This 'principle' of proportionate effort is often violated by many simulation tools. Investigators are often forced to accept a range of default parameters because they have no available data about particular aspects of an incident. Some automobile simulators require information about whether lights and wipers were working at the point of impact;

- *Regulatory guidelines can restrict the use of simulation tools.* Previous paragraphs have referred to the 'god's eye' view that is provided by some computer-based systems. These tools enable investigators to integrate data from diverse sources, many of which could not have been accessed by operators during an incident. The resulting animations can have an insidious effect. With the benefit of hindsight and with access to these additional data sources they can used to suggest that operators should have been better prepared for an incident. Reasonable concerns over this mis-use of technology have persuaded some regulators to restrict the use of computer-based simulations [423]. Typically, these only apply to the public dissemination of any resulting animations or models rather than to their use in the reconstruction of an incident.

These objections must be balanced against the benefits of computer-based simulations. These can be summarised as follows:

- *Simulations enable investigators to reconstruct the environment in which an incident occurs.* As we have seen, a range of model-based and photorealistic techniques can be used to reconstruct the layout of a working environment. Investigators can then use interpolation and rendering software to move within those environments. This approach can be used to record the aftermath of an incident, for instance using QuicktimeVR techniques. It can also be used to recreate what the operators might have seen from a number of different locations. The same approaches can also be used to reconstruct particular items that were involved in an incident. A common strength of all these visualisation techniques is that enable users to survey sites that may be too hazardous or expensive for them to visit on subsequent occasions.

- *Simulations enable investigators to replay events in real and virtual time.* The opening sections of this chapter stressed the role of reconstructions in establishing agreement over the course of an incident. Computer-based animations offer considerable flexibility in the way that investigators can play and replay particular events. Software support can be used to alter the real-time of key failures. The sequence in which they occur can also be reviewed and amended. Although there may be initial costs in terms of the time taken to learn how to exploit this authoring software, these are more that off-set by the flexibility of the resulting simulations. In particular, computer-based simulation tools can significantly redice the complexity associated with the maintenance of paper-based documents capture many hundreds of events;

- *Simulations enable investigators to integrate multiple data sources.* Primary and secondary investigations are intended to secure incident data from many diverse sources. Computer-based simulations provide means of integrating this information into hybrid reconstructions. There are many existing problems. For instance, it can be difficult to integrate continuous and discrete data. It can also be difficult to determine the best means of exploiting probabilistic information. Current research continues to explore means of representing operator intervention. In spite of these caveats, computer-based simulations arguably provide the greatest hope of unifying the increasing mass of information that can be obtained in the aftermath of many safety-critical incidents;

- *Simulations enable investigators to explore 'subjunctive' behaviours.* There is an increasing recognition that advanced mathematical techniques can be used to model effects that are difficult or impossible to assess using other approaches. For instance, the effects of muscle tension during traffic accidents can be tested using dummies or cadavers. In the former case, there are considerable recalibration problems associated with the linkages between simulated muscle groups. In the latter case, additional support structures must be used to maintain posture prior to the simulated incident. Alternatively, computer-based models can be derived to provide a low-cost means of simulating the consequences of crashes again and again and again. The results of these studies can be validated using more conventional techniques. However, the costs of crashing full-scale replicas would prohibit the range and scope of tests that can easily be conducted using software environments. An important benefit of these approaches is that support subjunctive simulation. In other words, the low-costs associated with running a trial can encourage investigators to consider a wide range of 'what if' scenarios. This, in turn, encourages the generalisation that was emphasised in the opening sections of this chapter. It is possible to llok beyond the specific events of a near-miss incident to examine potential ways in which such failures might have had more serious consequences.

The use of computer-based reconstruction in incident and accident investigation has had a number of notable successes [9, 211]. It is important to emphasise, however, that many questions remain to be answered about the pragmatic application of these techniques. The development of reconstruction tools to support incident reconstruction lags behind other applications of this technology. The US Aviation Safety Research Act of 1988 provides an insight into the reasons for this lack of investment. This act charged the FAA to undertake "...a research program to develop dynamic simulation models of the air traffic control (ATC) system which will provide analytical technology for predicting airport and ATC safety and capacity problems, and for evaluating planned research projects". The US National Simulation Capability Program was established in response to this congressional mandate. Its goals and objectives reflect the intentions of the Act. Simulations are intended to support training and assist in the development of new systems. The Act arguably neglects the role that simulations can play in understanding the causes of past failures.

Chapter 9

Modelling Notations

The previous chapter has introduced the growing number of computer-based simulation tools that can be used to reconstruct and replay the events leading to failure. As we have seen, however, these tools can be costly both to purchase and to apply to particular incidents. There is also the significant danger that they may bias investigators towards particular conclusions. For example, it is far easier to simulate the direct physical failure of component hardware than it is to model the managerial or regulatory failures that created the latent conditions for an incident. In consequence, this chapter introduces a number of notations that can be used to reconstruct the events leading to adverse occurrences. These techniques range from relatively 'intuitive' extensions to text-based time-lines through more complex graphical notations, such as Petri Nets, to mathematical logic. The intention is not to advocate a particular technique but to illustrate the costs and benefits of each approach. The final sections build on this analysis by presenting a list of requirements to be satisfied by any abstract notation for incident reconstruction.

9.1 Reconstruction Techniques

As we have seen, incidents may take many days, weeks or even years to develop [699]. As a result, a range of reconstruction techniques have been developed to help investigators represent and reason about the events that contribute to adverse occurrences. The following paragraphs briefly introduce a number of these approaches. These include graphical time-lines which provide a sketch of the events leading to an incident. We also consider the application of Fault trees to support the reconstruction of adverse occurrences. This diagrammatic techniques has been widely applied to support systems development and is, therefore, accessible to many engineers and investigators. Later sections also explore the use of Petri Nets reconstructions. This graphical notation is specifically intended to represent and reason about the complex temporal properties that characterise many incidents. A textual logic is then considered. This approach lacks the visual appeal of the graphical notations but has well-developed proof techniques that enable investigators to establish key properties of any reconstruction.

In order to illustrate the application of these different reconstruction techniques, the following sections analyse an incident involving the rupture of a natural gas distribution pipeline [588]. A 2-inch-diameter steel gas service line had been exposed during the excavation that was intended to help the with the removal of a 8,000 gallon buried fuel tank. The exposed pipeline separated at a compression coupling about 5 feet from the wall of a retirement home in Allentown, Pennsylvania. The escaping gas flowed underground, passed through openings in the building foundation. It then migrated to other floors in the retirement home before it exploded. The Allentown incident resulted in one fatality. The consequence criteria that were introduced in the opening chapters could, therefore, be used to argue that this is an accident and not a 'near-miss' incident. However, pragmatic and theoretical justifications support the decision to use this case study. The pragmatic explanation is that the subsequent National Transportation Safety Board (NTSB) report provides detailed information about the secondary investigation of this explosion. This provides public access to the sorts of details that often remain confidential within many commercial reporting systems. The theoretical justification for using this incident is that its consequences could have been very much worse. Many of the elderly residents of the retirement home were not in the building at the time of the explosion. A final motivation for using this incident is that it represents one of a number of similar incidents, the causes of which had arguably not been properly addressed by the pipeline and construction industries or their regulators.

9.1.1 Graphical Time Lines

Time-lines are one of the simplest means of representing the flow of events during major accidents. They simply translate the events on the text-based time-lines, which have been presented in previous paragraphs, onto a horizontal or vertical axis. Each event is mapped to a point on a line which stretches from the earliest to the lastest moment that is considered to be important to the analysis. For example, Figure 9.1 examines the regulatory environment in which the Allentown incident took place. In particular, this diagram provides a high-level overview of Appendix B of the NTSB's report. This natural language account is over twenty pages long. The graphical time-line does not replace the more detailed prose, however, it does provide an overview of the information that it contains. It focuses on the way in which the NTSB and groups within the Department of Transportation (DOT), in particular the Office of Pipeline Safety (OPS) within the Research and Special Programs Administration (RSPA), responded to previous incidents. In particular, it looks at the way in which recommendations to introduce Excess Flow Valves (EFVs) had limited uptake in the industry. The NTSB report argues that these devices could have mitigated the consequences of the Allentown incident. As can be seen, EFV devices were initially pioneered in the late 1960s. Incidents in 1968 and in 1972 had led the NTSB to recommend that the OPS develop standards for the use of protection devices such as EFVs. As a result, OPS recommend the installation of Excess Flow Valves (EFVs) in all new gas service lines and lines undergoing repair in 1974. Incidents continued to occur and later the same year, a Department of Transportation report recommended that EFVs be extended to customer lines. In 1976 the NTSB recommended their use in commercial premises. However, the Office of Pipeline Safety argued that the results of tests on these devices had proved to be inconclusive.

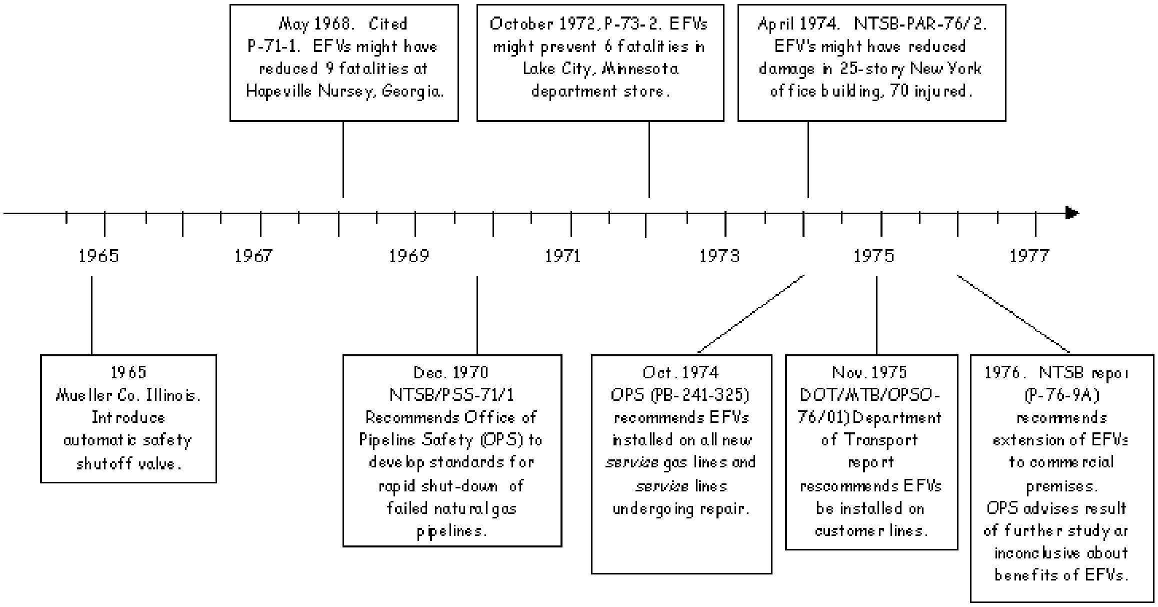
May 1968. Cited P-71-1. EFVs might have reduced 9 fatalities at Hapeville Nursey, Georgia.
October 1972, P-73-2. EFVs might prevent 6 fatalities in Lake City, Minnesota department store.
April 1974. NTSB-PAR-76/2. EFV's might have reduced damage in 25-story New York office building, 70 injured.
1965
1967
1969
1971
1973
1975
1977
1965 Mueller Co. Illinois. Introduce automatic safety shutoff valve.
Dec. 1970 NTSB/PSS-71/1 Recommends Office of Pipeline Safety (OPS) to develop standards for rapid shut-down of failed natural gas pipelines.
Oct. 1974 OPS (PB-241-325) recommends EFVs installed on all new *service* gas lines and *service* lines undergoing repair.
Nov. 1975 DOT/MTB/OPSO-76/01) Department of Transport report rescommends EFVs be installed on customer lines.
1976. NTSB repo (P-76-9A) recommends extension of EFVs to commercial premises. OPS advises result of further study ar inconclusive about benefits of EFVs.

Figure 9.1: Graphical Time-line Showing Initial Regulatory Background.

The spatial layout of this time-line is not simply used to indicate the flow of events over time. In

Figure 9.1 previous incidents are grouped above the line. The regulatory responses to those events are grouped below the line. This format is intended to show the impact the previous failures had on wider aspects of safety management within the pipeline and construction industries. Figure 9.1.1 builds on this approach by extending the time-line closer to the Allentown incident. As can be seen, a number of further incidents occurred that either had similar causes to the Allentown incident, such as the Green County incident, or in which the NTSB again recommended the use of EFVs. Again, the labels below the time-line are used to chart the progress of regulatory studies and recommendations about the use of EFVs. This diagram shows the NTSB's continuing support for the wider introduction of these devices, for example in the 1981 study of 14 previous accidents. It also illustrates concerns about the reliability and utility of these devices within the Department of Transportation. These concerns lead to a study by the Gas Research Institute in 1985. The NTSB concludes that this report is seriously flawed in 1987-88 in that it under-estimates the utility of EPVs.

Figure 9.1.1 further extends the previous two time-lines beyond the Allentown explosion. It illustrates the continuing debate between the regulator, the Department of Transportation, and the investigatory agency, the NTSB. Following the Allentown explosion, a group of seventeen congressional representatives signed a letter that was sent to the Department of Transportation. This criticised the lack of progress that had been made in improving pipeline safety. However, the Office of Pipeline Safety still deferred any final ruling on the widespread introduction of EFVs.

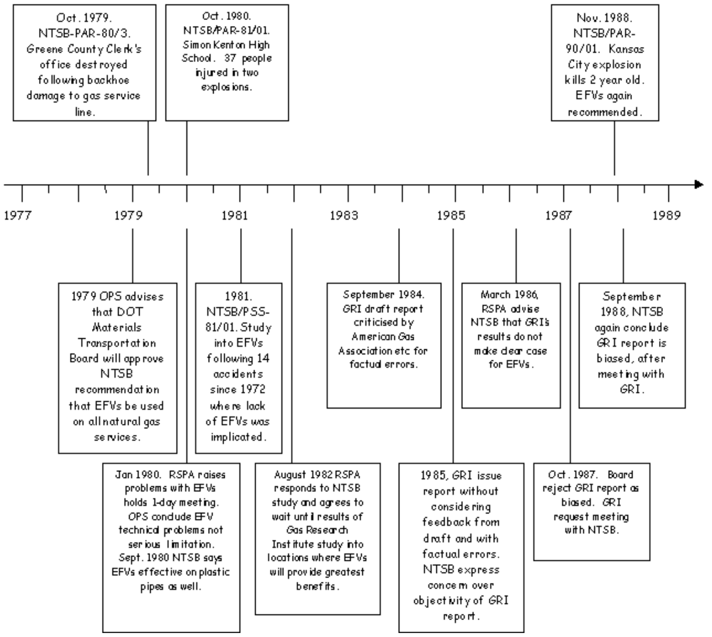
Oct. 1979. NTSB-PAR-80/3. Greene County Clerk's office destroyed following backhoe damage to gas service line.
Oct. 1980. NTSB/PAR-81/01. Simon Kenton High School. 37 people injured in two explosions.
Nov. 1988. NTSB/PAR-90/01. Kansas City explosion kills 2 year old. EFVs again recommended.
1977
1979
1981
1983
1985
1987
1989
1979 OPS advises that DOT Materials Transportation Board will approve NTSB recommendation that EFVs be used on all natural gas services.
1981. NTSB/PSS-81/01. Study into EFVs following 14 accidents since 1972 where lack of EFVs was implicated.
September 1984. GRI draft report criticised by American Gas Association etc for factual errors.
March 1986, RSPA advise NTSB that GRI's results do not make clear case for EFVs.
September 1988, NTSB again conclude GRI report is biased, after meeting with GRI.
Jan 1980. RSPA raises problems with EFVs holds 1-day meeting. OPS conclude EFV technical problems not serious limitation. Sept. 1980 NTSB says EFVs effective on plastic pipes as well.
August 1982 RSPA responds to NTSB study and agrees to wait until results of Gas Research Institute study into locations where EFVs will provide greatest benefits.
1985, GRI issue report without considering feedback from draft and with factual errors. NTSB express concern over objectivity of GRI report.
Oct. 1987. Board reject GRI report as biased. GRI request meeting with NTSB.

Figure 9.2: Graphical Time-line Showing Intermediate Regulatory Background.

As has been mentioned in previous paragraphs, the graphical time-lines provide a framework or overview of the events that contribute to an incident or accident. Each entry can be thought of as an index into the more detailed evidence that is gathered during a secondary investigation. It also follows that not all of this evidence may be shown on a graphical time-line. Reconstruction, typically, involves a process of abstraction that is implied by our use of terms such as 'overview' or 'model' in the previous paragraphs. Investigators must use their judgement to determine what is, and what is not, included in a time-line. For instance, the previous diagrams have not included the letter that Jim Hall, Chairman of the NTSB, sent on the 28th September 1995 to the administrator of the Research and Special Programs Administration in the Department of Transport. The recipient of this letter was responsible for managing the Office of Pipeline Safety. Chairman Hall's letter expressed disappointment at the RSPA's response to House and Senate committees when they failed

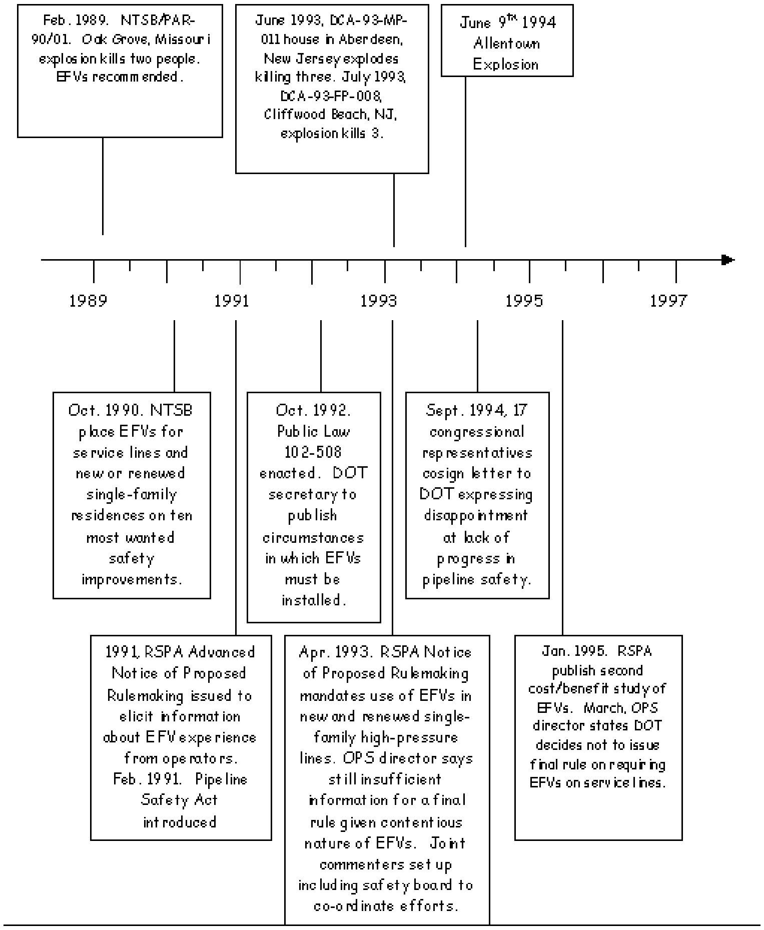
Feb. 1989. NTSB/PAR-90/01. Oak Grove, Missouri explosion kills two people. EFVs recommended.
June 1993, DCA-93-MP-011 house in Aberdeen, New Jersey explodes killing three. July 1993, DCA-93-FP-008, Cliffwood Beach, NJ, explosion kills 3.
June 9th 1994 Allentown Explosion
1989
1991
1993
1995
1997
Oct. 1990. NTSB place EFVs for service lines and new or renewed single-family residences on ten most wanted safety improvements.
Oct. 1992. Public Law 102-508 enacted. DOT secretary to publish circumstances in which EFVs must be installed.
Sept. 1994, 17 congressional representatives cosign letter to DOT expressing disappointment at lack of progress in pipeline safety.
1991, RSPA Advanced Notice of Proposed Rulemaking issued to elicit information about EFV experience from operators. Feb. 1991. Pipeline Safety Act introduced
Apr. 1993. RSPA Notice of Proposed Rulemaking mandates use of EFVs in new and renewed single-family high-pressure lines. OPS director says still insufficient information for a final rule given contentious nature of EFVs. Joint commenters set up including safety board to co-ordinate efforts.
Jan. 1995. RSPA publish second cost/benefit study of EFVs. March, OPS director states DOT decides not to issue final rule on requiring EFVs on service lines.

Figure 9.3: Graphical Time-line Showing Immediate Regulatory Background.

to identify any circumstances that might mandate the use of EFVs. The Chairman of the NTSB continued:

> "The Safety Board is extremely disappointed in your decision. For more than 20 years, RSPA has failed to objectively assess the benefits of EFVs, and we believe RSPA has again lost an excellent opportunity to provide increased safety for gas customers and the public... In our investigations of distribution pipeline accidents, the Safety Board continues to find strong evidence that supports requiring a means to rapidly shut off gas flow to failed pipe segments. While such a requirement would not prevent accidents, it would significantly reduce their consequences." [588]

The previous time-lines illustrate some of the production problems that limit the tractability of this approach. Initially, Figure 9.1 9.1.1 and 9.1.1 formed part of a single time-line. However, it proved to be impossible to reproduce this within the format and pagination of this book. As a result, the simple spatial relationship between layout and time had to be broken, in part, by splitting a single linear diagram into several different figures. Later paragraphs will show how hierarchical time-lines can be used to avoid this potential limitation.

The high-level time-lines shown in Figures 9.1, 9.1.1 and 9.1.1 illustrate many of the strengths of this reconstruction or modelling technique. The simple relationship between spatial locations on the diagrams and temporal locations during an incident has already been noted. The practical consequence of this is that analysts need minimal training to use these models. They can be used

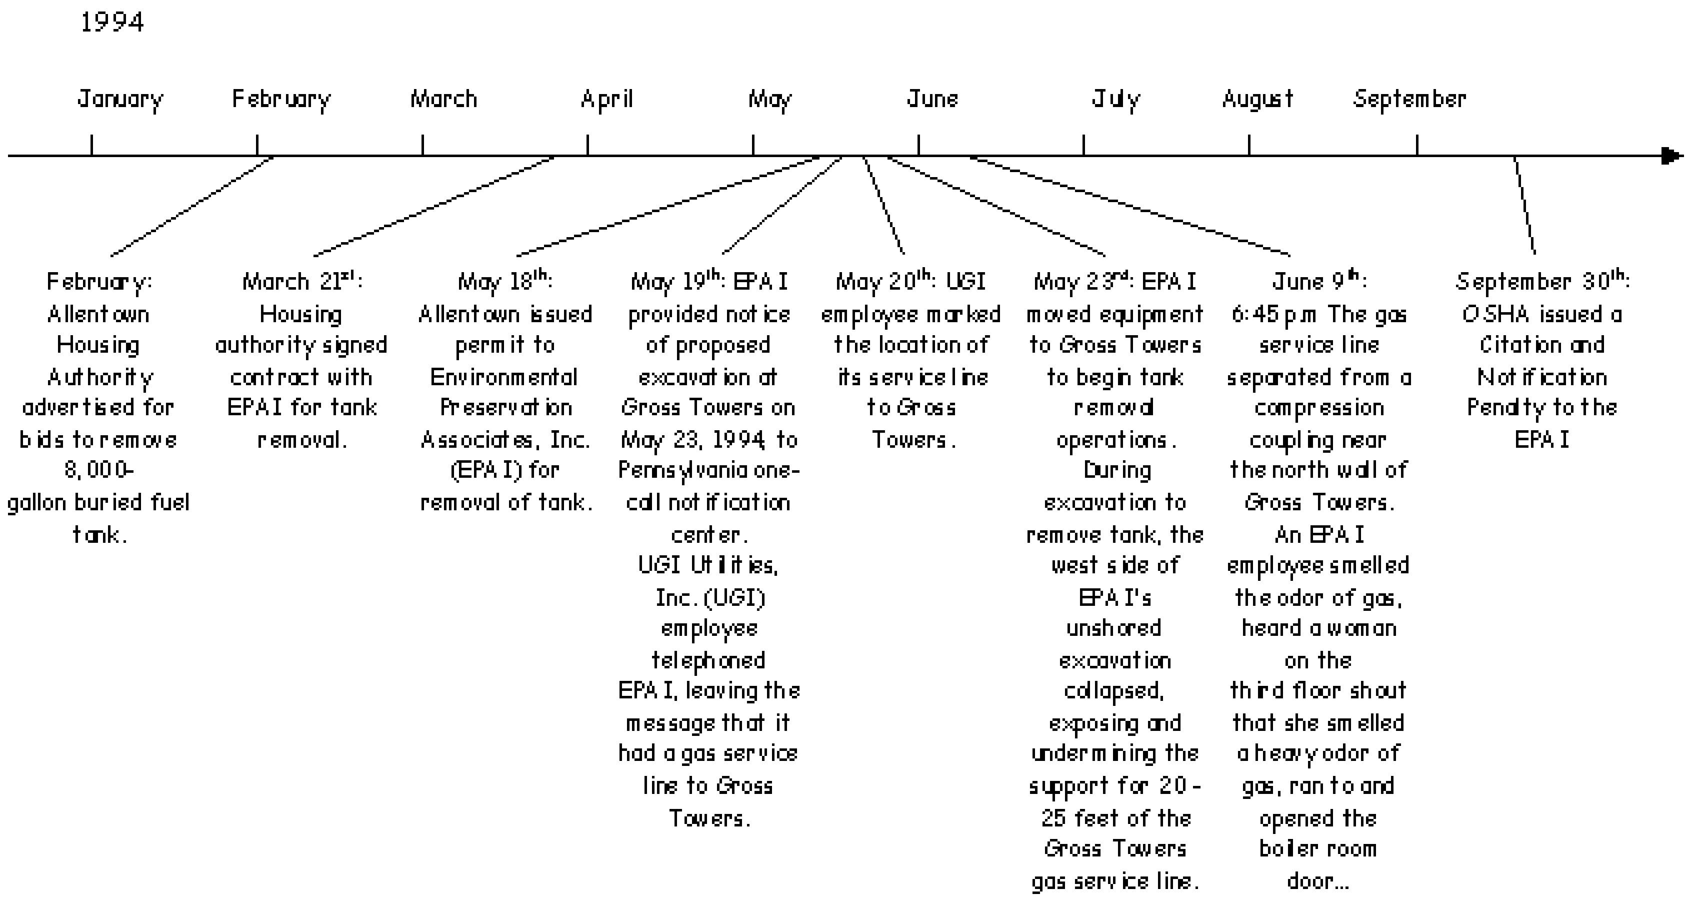

February: Allentown Housing Authority advertised for bids to remove 8,000-gallon buried fuel tank.

March 21st: Housing authority signed contract with EPAI for tank removal.

May 18th: Allentown issued permit to Environmental Preservation Associates, Inc. (EPA I) for removal of tank.

May 19th: EPA I provided notice of proposed excavation at Gross Towers on May 23, 1994 to Pennsylvania one-call notification center. UGI Utilities, Inc. (UGI) employee telephoned EPA I, leaving the message that it had a gas service line to Gross Towers.

May 20th: UGI employee marked the location of its service line to Gross Towers.

May 23rd: EPA I moved equipment to Gross Towers to begin tank removal operations. During excavation to remove tank, the west side of EPA I's unshored excavation collapsed, exposing and undermining the support for 20 - 25 feet of the Gross Towers gas service line.

June 9th: 6:45 p.m The gas service line separated from a compression coupling near the north wall of Gross Towers. An EPA I employee smelled the odor of gas, heard a woman on the third floor shout that she smelled a heavy odor of gas, ran to and opened the boiler room door...

September 30th: OSHA issued a Citation and Notification Penalty to the EPA I

Figure 9.4: Graphical Time-line of Events Surrounding the Allentown Explosion.

as a common medium of communication between the diverse disciplines involved in incident investigations. Figure 9.4 extends this approach by presenting a time-line for some of the events that relate more directly to the Allentown case study, rather than to a more general class of pipeline failures. Here we can see the strong visual appeal of this linear notation. Readers can easily gauge the intervals between events because there is a simple relationship between linear distance and the temporal intervals between events. In other words, standard units of distance are used to represent standard units of time. In Figure 9.4, this is used to indicate the interval between the date at which the Allentown Housing Association put the removal of the buried fuel tank out to tender and the date when Occupational Safety and Health Administration (OSHA) cited EPAI for a range of health and safety deficiencies. As can be seen, this diagram shows both the events leading to the gas line separation on the 9th June as well as events after the incident, such as the OSHA citation. This satisfies the reconstruction requirement that it should be possible to represent the consequent actions following any adverse occurrence. However, this graphical time-line illustrates events at an extremely high level of granularity. In contrast, Figure 9.5 shows how the same approach can be applied to the more detailed proximal events 'on the day of the incident' rather than the more distal causes shown in Figures 9.4, 9.1, 9.1.1. Unfortunately, Figure 9.5 illustrates further limitations with this reconstruction technique. In particular, it is necessary to position all events at some time during the incident. This is not always possible. For instance, the it was never possible to determine the exact time at which the foreman asked his team to trace the gas line back towards Utica street so that they could shut-off the gas valve. As a result this is labelled as occurring at 18:?? in Figure 9.5 and no connection can be made to the intervals illustrated on the time-line.

It is also possible to see an 'uneven' distribution of events over time in the clustering between 18:40 and 18:50. Nothing significant is shown to happen between the EPAI foreman's warning to the Housing Association Official that the gas line needed to be supported and the arrival of the backhoe at Gross Towers. Conversely, a large number of critical events take place in the interval between the moment when the backhoe was driven across the buried section of pipe and the moment when the foreman rang UGI to inform them that they had definitely hit the gas line. The concentration of critical events crams many different annotations into a small area of the line. This reduces the tractability of the resulting time-line.

This uneven distribution of events over time partially explains the decision to use two different scales in Figures 9.4 and 9.5. The former divides the line into months while the latter uses hours. If the same hour-based scale were used then the tendering process in February would have to be drawn many meters away from the events in May or June. Most of this line would have no significant annotations until the contract was signed in March. Although our use of different temporal scales helps to avoid this problem, it also introduce further concerns. In particular, investigators now have to maintain multiple diagrams of the same incident. Extensive cross references have to made in order to get a coherent overview of these different aspects of the same reconstruction. Figure 9.6 addresses this concern by using different axes to link the previous two graphical time-lines. This represents one of the hierarchical approaches mentioned in previous paragraphs . The higher-level intervals that are represented on one axis can be broken down into more fine grained intervals that are represented on an orthogonal axis. Unfortunately, this approach introduces further problems. In particular, it can be argued that Figure 9.6 destroys the simple, linear relationship between space and time that is claimed to be the key strength of the time-line notation. The following sections, therefore, describe a number of further reconstruction techniques that can be used to address these limitations of graphical time-lines.

9.1.2 Fault Trees

Fault trees provide an alternative means of reconstructing the events that contribute to incidents and accidents [502, 407]. This notation provides a simple graphical syntax based around circuit diagrams. Figure 9.7 presents a brief overview of the syntactic elements that support this approach. These elements are used to construct a diagram that connects basic events to higher-level faults. AND gates can be used to represent that a particular fault or intermediate event is caused by the combination of two or more basic events. Similarly, OR gates can be used to represent that a par-

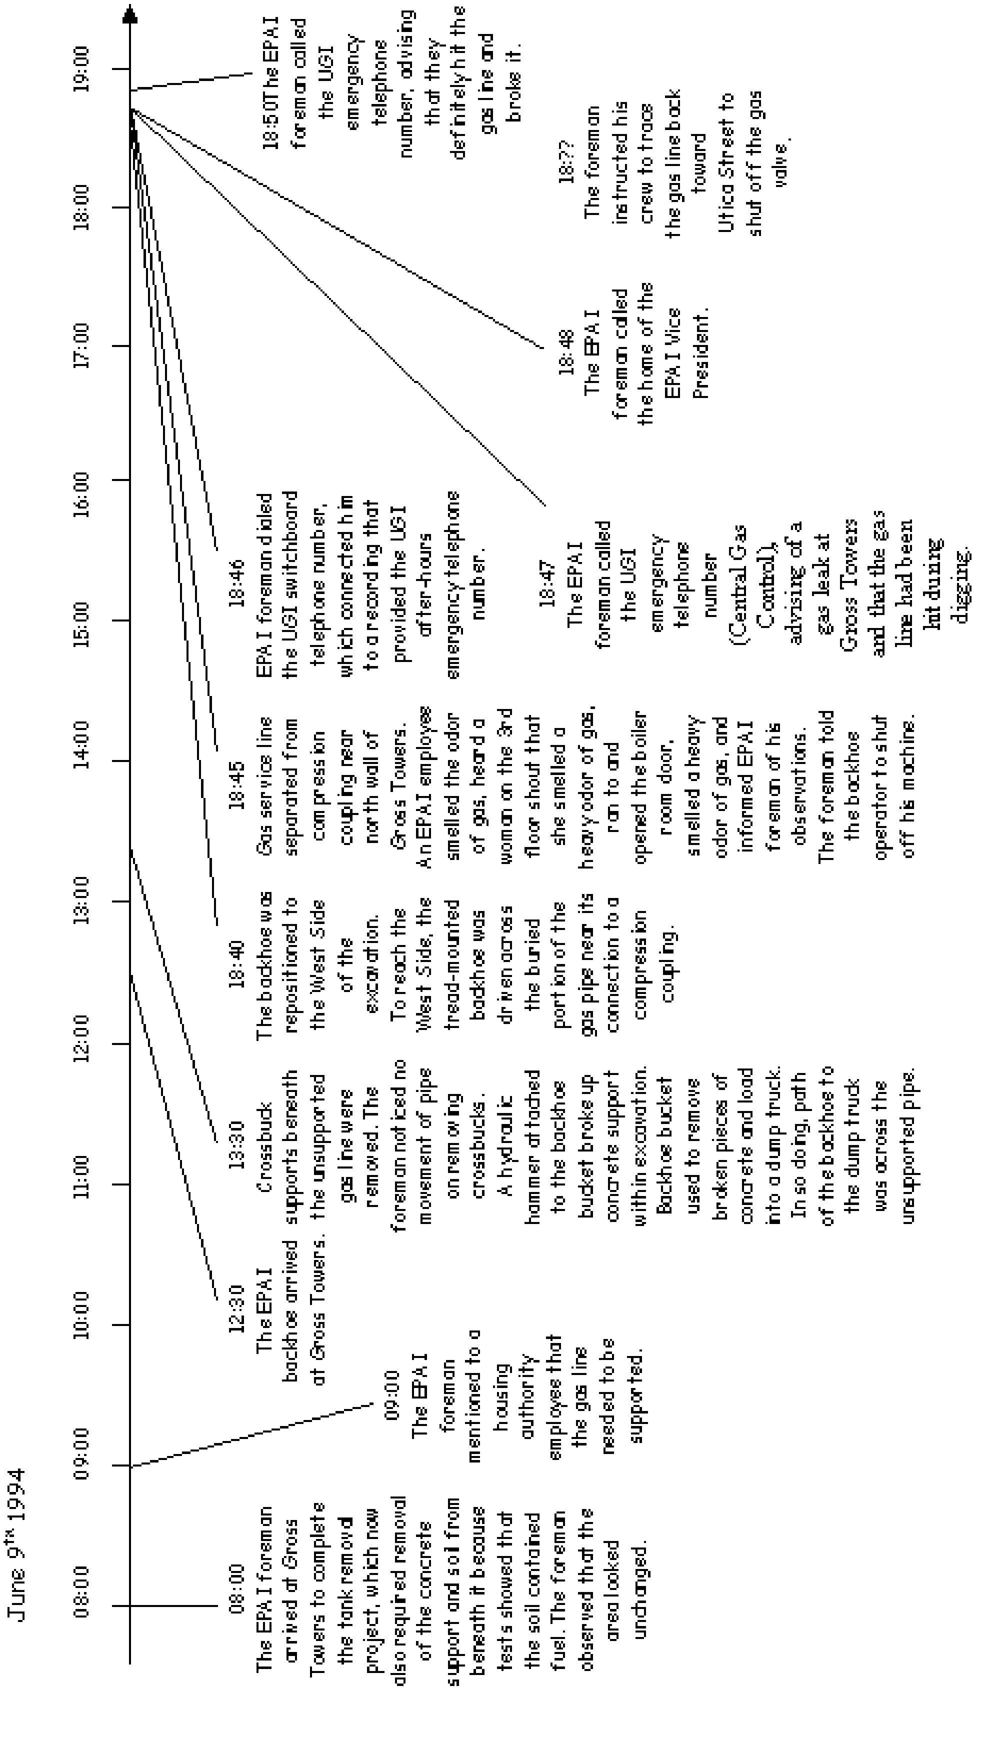
June 9th 1994
08:00
09:00
10:00
11:00
12:00
13:00
14:00
15:00
16:00
17:00
18:00
19:00
08:00 The EPA I foreman arrived at Gross Towers to complete the tank removal project, which now also required removal of the concrete support and soil from beneath it because tests showed that the soil contained fuel. The foreman observed that the area looked unchanged.
09:00 The EPA I foreman mentioned to a housing authority employee that the gas line needed to be supported.
12:30 The EPA I backhoe arrived at Gross Towers.
13:30 Crossbuck supports beneath the unsupported gas line were removed. The foreman noticed no movement of pipe on removing crossbucks. A hydraulic hammer attached to the backhoe bucket broke up concrete support within excavation. Backhoe bucket used to remove broken pieces of concrete and load into a dump truck. In so doing, path of the backhoe to the dump truck was across the unsupported pipe.
18:40 The backhoe was repositioned to the West Side of the excavation. To reach the West Side, the tread-mounted backhoe was driven across the buried portion of the gas pipe near its connection to a compression coupling.
18:45 Gas service line separated from compression coupling near north wall of Gross Towers. An EPA I employee smelled the odor of gas, heard a woman on the 3rd floor shout that she smelled a heavy odor of gas, ran to and opened the boiler room door, smelled a heavy odor of gas, and informed EPA I foreman of his observations. The foreman told the backhoe operator to shut off his machine.
18:46 EPA I foreman dialed the UGI switchboard telephone number, which connected him to a recording that provided the UGI after-hours emergency telephone number.
18:47 The EPA I foreman called the UGI emergency telephone number (Central Gas Control), advising of a gas leak at Gross Towers and that the gas line had been hit during digging.
18:48 The EPA I foreman called the home of the EPA I Vice President.
18:?? The foreman instructed his crew to trace the gas line back toward Utica Street to shut off the gas valve.
18:50 The EPA I foreman called the UGI emergency telephone number, advising that they definitely hit the gas line and broke it.

Figure 9.5: Graphical Time-line of the Allentown Explosion.

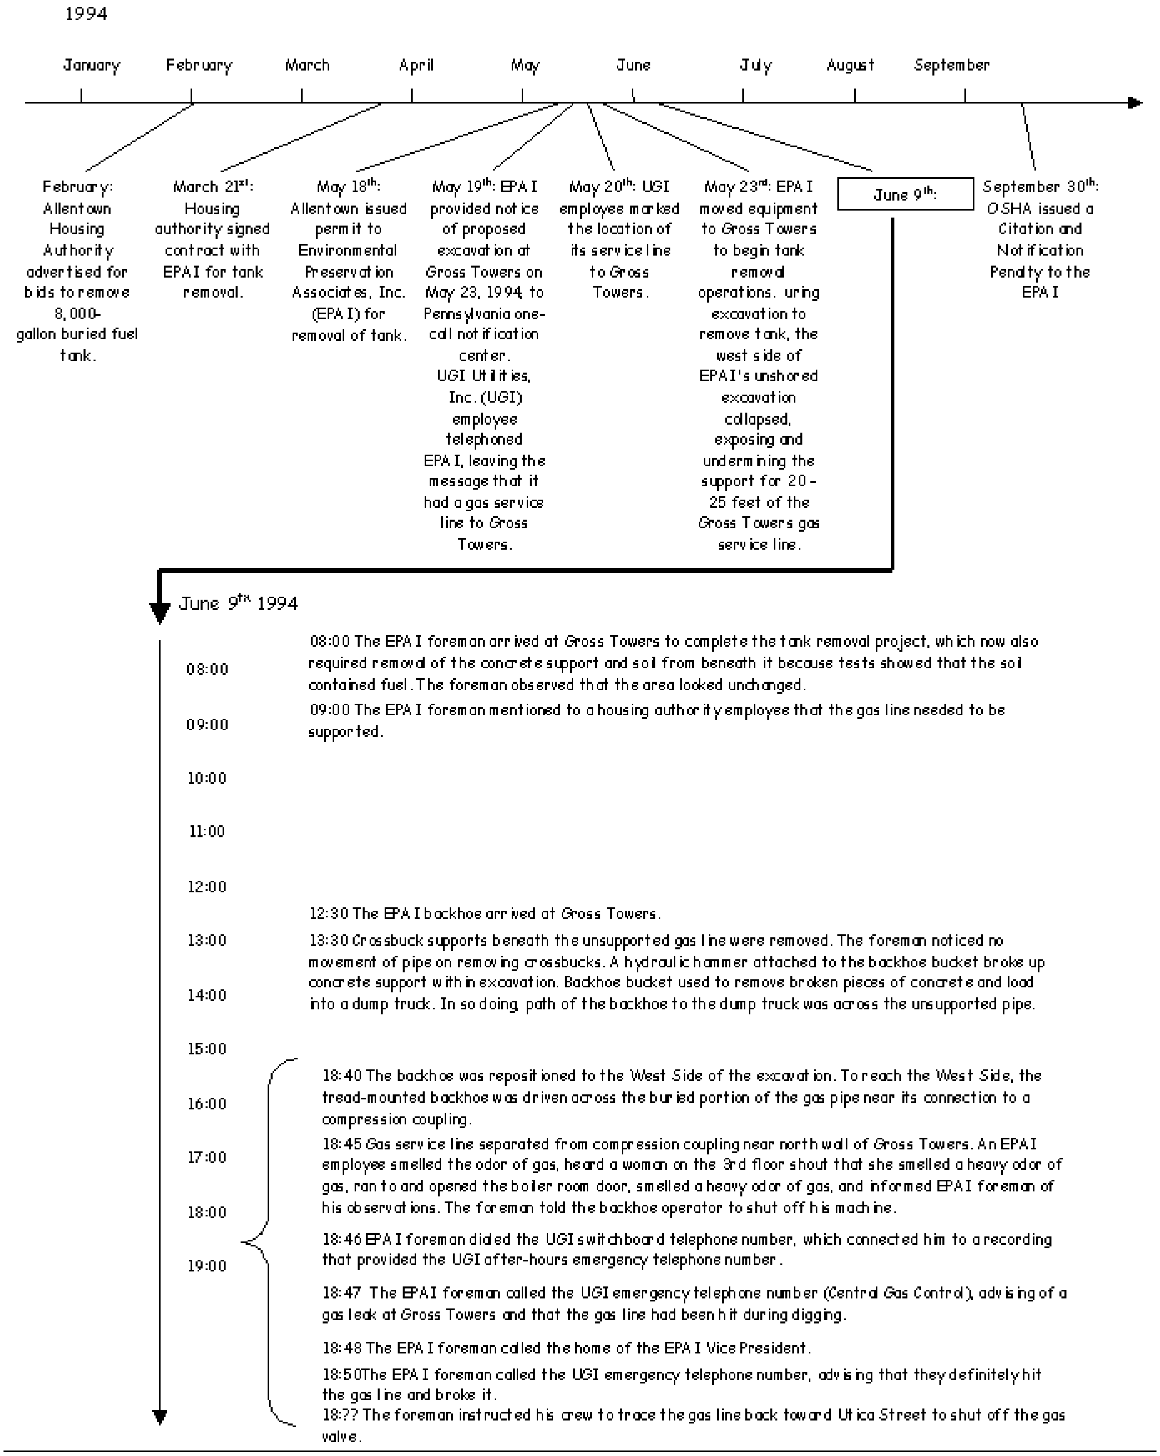
1994
January
February
March
April
May
June
July
August
September
February: Allentown Housing Authority advertised for bids to remove 8,000-gallon buried fuel tank.
March 21st: Housing authority signed contract with EPAI for tank removal.
May 18th: Allentown issued permit to Environmental Preservation Associates, Inc. (EPAI) for removal of tank.
May 19th: EPAI provided notice of proposed excavation at Gross Towers on May 23, 1994 to Pennsylvania one-call notification center. UGI Utilities, Inc. (UGI) employee telephoned EPAI, leaving the message that it had a gas service line to Gross Towers.
May 20th: UGI employee marked the location of its service line to Gross Towers.
May 23rd: EPAI moved equipment to Gross Towers to begin tank removal operations. uring excavation to remove tank, the west side of EPAI's unshored excavation collapsed, exposing and undermining the support for 20-25 feet of the Gross Towers gas service line.
June 9th:
September 30th: OSHA issued a Citation and Notification Penalty to the EPAI
June 9tx 1994
08:00
09:00
10:00
11:00
12:00
13:00
14:00
15:00
16:00
17:00
18:00
19:00
08:00 The EPAI foreman arrived at Gross Towers to complete the tank removal project, which now also required removal of the concrete support and soil from beneath it because tests showed that the soil contained fuel. The foreman observed that the area looked unchanged.
09:00 The EPAI foreman mentioned to a housing authority employee that the gas line needed to be supported.
12:30 The EPAI backhoe arrived at Gross Towers.
13:30 Crossbuck supports beneath the unsupported gas line were removed. The foreman noticed no movement of pipe on removing crossbucks. A hydraulic hammer attached to the backhoe bucket broke up concrete support within excavation. Backhoe bucket used to remove broken pieces of concrete and load into a dump truck. In so doing, path of the backhoe to the dump truck was across the unsupported pipe.
18:40 The backhoe was repositioned to the West Side of the excavation. To reach the West Side, the tread-mounted backhoe was driven across the buried portion of the gas pipe near its connection to a compression coupling.
18:45 Gas service line separated from compression coupling near north wall of Gross Towers. An EPAI employee smelled the odor of gas, heard a woman on the 3rd floor shout that she smelled a heavy odor of gas, ran to and opened the boiler room door, smelled a heavy odor of gas, and informed EPAI foreman of his observations. The foreman told the backhoe operator to shut off his machine.
18:46 EPAI foreman dialed the UGI switchboard telephone number, which connected him to a recording that provided the UGI after-hours emergency telephone number.
18:47 The EPAI foreman called the UGI emergency telephone number (Central Gas Control), advising of a gas leak at Gross Towers and that the gas line had been hit during digging.
18:48 The EPAI foreman called the home of the EPAI Vice President.
18:50The EPAI foreman called the UGI emergency telephone number, advising that they definitely hit the gas line and broke it.
18:?? The foreman instructed his crew to trace the gas line back toward Utica Street to shut off the gas valve.

Figure 9.6: Two-Axis Time-line of the Allentown Explosion.

ticular fault or intermediate event is caused by some subset of the more basic events that are linked to it. An exclusive-OR gate can be used to further restrict this representation so that a particular fault or intermediate event is caused by one of a number of more basic events. Andrews and Moss [27] provide a more detailed introduction to the fault tree notation. However, the following paragraphs will introduce the basic concepts as they are used. Fault trees are, typically, used to

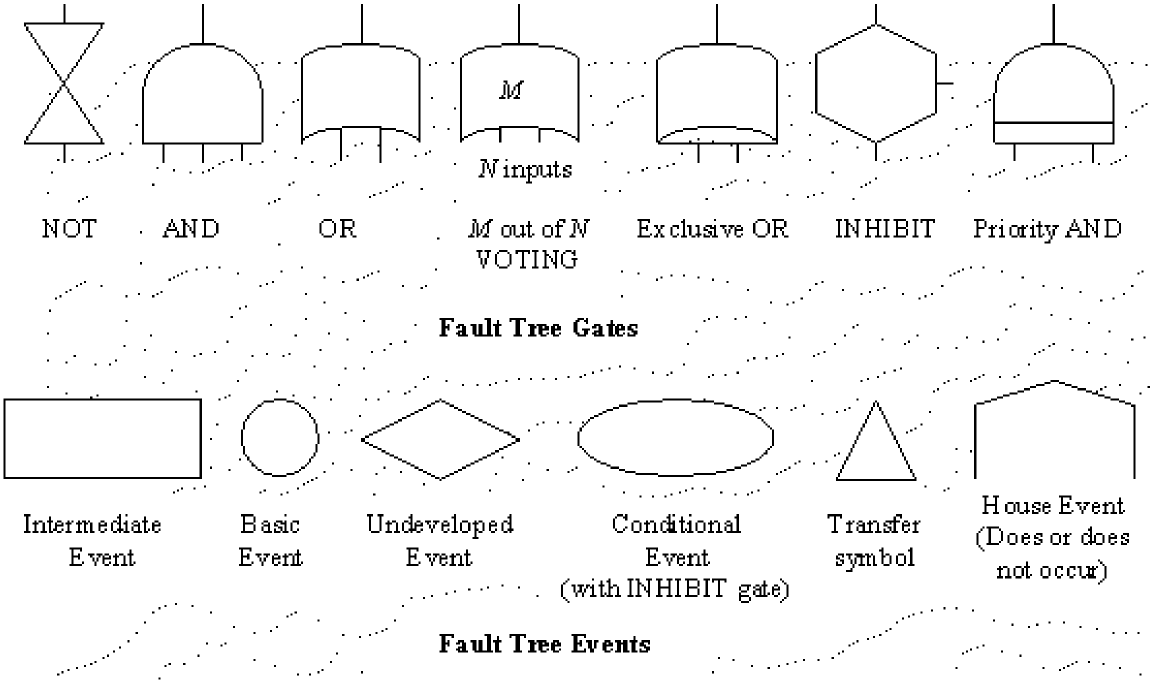

Figure 9.7: Fault tree components.

analyse potential errors in a design. This is illustrated by the simplified tree shown in figure 9.8. An operator injury occurs if three conditions are met. The protective guard must fail and a command to initiate the press must be given and the operator's hand must be under the protective guard. As can be seen, the conjunction between all of these three conditions is denoted by the graphical symbol that represents an AND operation within a circuit diagram. The left hand sub-branch of Figure 9.8 shows two ways in which the guard can fail. There may be a physical obstruction that prevents the guard from closing or an electrical failure of the guard motor may occur while it is still in the open position. Here the disjunction between these two conditions is denoted by the graphical symbol that represents an OR operation within a circuit diagram. There are numerous design benefit for this, typical, application of fault trees. For instance, they can be used to identify what is known as the minimal cut set. In order to explain this concept it is first necessary to explain that a basic event is one which cannot be decomposed any further. In figure 9.8 'Physical obstruction blocks guard at open' is a basic event. In contrast, 'guard fails' is an intermediate or higher level event. A minimal cut set is defined to be the smallest possible conjunction of events in which if any basic event is removed then the top condition will not occur [27]. For our example, there are two minimal cut sets. Operator injury will occur if:

```
Physical obstruction blocks guard AND
Pressing initiated AND
Operator's hand is under guard

OR
```

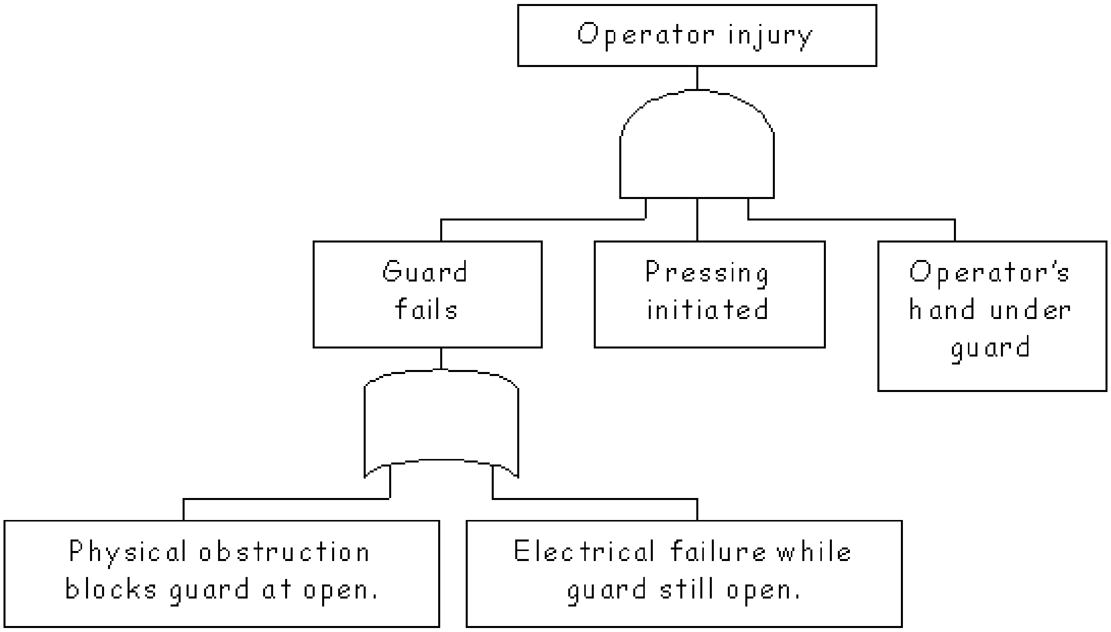
Operator injury
Guard fails
Pressing initiated
Operator's hand under guard
Physical obstruction blocks guard at open.
Electrical failure while guard still open.

Figure 9.8: A Simple Fault Tree for Design.

```
Electrical failure while guard at open AND
Pressing initiated AND
Operator's hand is under guard
```

The importance of a minimal cut set is that it can be used to identify where to focus finite development resources. If there is a basic event that is common to all minimal cut sets and it is possible to prevent that event from occurring then, by definition, the top event cannot also occur. This assumes that the fault tree accurately reflects all of the possible ways in which an adverse occurrence can take place. Conversely, if it is only possible to prevent basic events that occur in some proportion of the minimal cut sets then there will continue to be other ways in which the incident may occur. Extensions of this basic approach can also be used to analyse the probability of a top level event if designers know the probability of the basic events that contribute to it. For instance, if observations of previous operations suggest that a physical obstruction blocks the guard once every hundred days then we assign a probability of it failing in the next day of 0.01. Similarly if observations suggest an electrical failure once every 1,000 days then the probability would be 0.001. The probability of the disjunction of an electrical failure, shown as event A, or of an obstruction, shown as event B, is derived by applying the following formula. The final term accounts for the situation in which both the electrical failure and the physical obstruction occur together.

$$Pr(A\ or\ B) = Pr(A) + Pr(B) - Pr(A\ and\ B) \tag{9.1}$$

If these events were mutually exclusive, in other words the physical obstruction and the electrical failure could not occur together, then this term could be omitted. In similar fashion, the probability of a conjunction can, most simply, be given as the product of the probabilities of its child events. There are, however, a number of technical and practical limitations. For example, it can be difficult to obtain reliable statistical data to validate the probabilities that are included in the tree. There are also a number of limitating assumptions, such as event independence. If these assumptions are violated then more complex mathematical procedures must be used to calculate conditional probabilities [27].

As mentioned, fault trees have traditionally been used to support the design of safety-critical systems. This notation can, however, offer considerable benefits as a means of supporting the reconstruction of adverse occurrences. The leaves of the tree represent the initial causes of an incident [485]. Basic events can be used to represent the underlying failures that lead to an accident [361]. Logic gates can be used to represent the ways in which those causes combine. For example, the combination of operator mistakes, hardware/software failures and managerial problems might be represented using an AND gate. Conversely, a lack of evidence about user behaviour or system performance might be represented using an OR gate. For instance, Figure 9.9 uses a fault tree to reconstruct part of the NTSB case study:

> "By reducing the soils capacity to restrain the movement of the pipe and by exerting forces on the service line that resulted in excessive longitudinal stress, the excavator caused the line to separate at a compression coupling. "
> "The gas company lost the opportunity to preserve the integrity of the service line because its procedures did not require a review of any unusual excavation near a gas service line that might damage the line and threaten public safety."
> "The likely reason the fire inspectors did not notify the gas company that its service line was damaged was because the inspectors did not understand the importance of notifying operators so the effects on a facility could be assessed by the operators and necessary action taken." [588]

As can be seen, the subtree on the right of Figure 9.9 represents the conjunction of events that are identified as causes for the line separation at the compression coupling: the soils capacity to restrain the movement of the pipe was reduced and undue forces were exerted on the line and gas company procedures did not require a thorough review of unusual excavations. Had any one of these events not taken place then the incident would not have occurred in the manner described by the NTSB. The counterfactual reasoning in the previous sentence illustrates the important point that the elements of a minimal cut-set within an accident fault tree are root causes of the ultimate failure that (paradoxically) is at the root, or top, of the entire tree structure. The events that contribute to the line separation, labelled Conclusion 3 and 6, and the failure of the fire inspectors to notify the company, labelled Conclusion 7, are all members of the minimal cut set and are, therefore, root causes of the gas explosion.

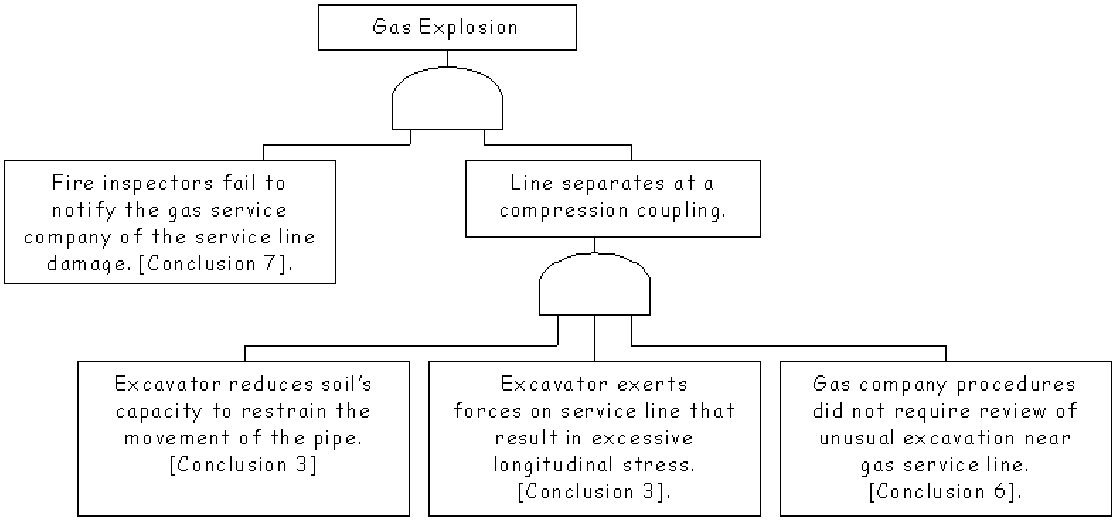
Gas Explosion
Fire inspectors fail to notify the gas service company of the service line damage. [Conclusion 7].
Line separates at a compression coupling.
Excavator reduces soil's capacity to restrain the movement of the pipe. [Conclusion 3]
Excavator exerts forces on service line that result in excessive longitudinal stress. [Conclusion 3].
Gas company procedures did not require review of unusual excavation near gas service line. [Conclusion 6].

Figure 9.9: Simplified Fault Tree Representing Part of the Allentown Incident.

The previous paragraph has shown how fault trees can be used to represent the root causes that are identified by counter factual reasoning. Unfortunately, this raises a number of practical and theoretical problems. As we have seen, our counter factual reasoning relies on that fact that the intermediate events described by an AND gate will only occur if all of its inputs are true. The gas explosion in Figure 9.9 would not have occurred if any of the four basic events were prevented from happening. This is an extremely strong assumption. How confident can we be that an explosion would actually have been avoided if the Fire inspectors had intervened? It is difficult to be certain that an incident would have been avoided under such circumstances. A number of complex reasoning techniques, based around modal logic, can be used to address this apparent limitation [470]. It is also possible to recruit additional forms of secondary investigation to increase our confidence in the elements of a reconstruction. In the previous example, this could involve further studies of the interaction between Fire inspectors and gas service companies. These studies might demonstrate that inspectors routinely intervene to prevent similar incidents from occurring. However, if there was strong evidence that such interventions have not been effective in avoiding gas explosions then the tree must be redrawn.

Immediate Causes

Figure 9.9 provides a high-level overview of some of the causes that led to the Allentown explosion. However, such abstract fault trees provide few insights into the more detailed patterns of events that contribute to major incidents. For example, Figure 9.9 abstracts away from the particular way in which the excavators' actions led to undue forces being exerted on the exposed gas line. Similarly, it does not identify the contextual or motivating factors that prevented the fire inspectors from notifying the damage that they observed to the gas company's line. Figure 9.10, therefore, shows how a fault tree can be used to provide a more detailed overview of the events leading to an adverse occurrence. This diagram is significantly more complex than its predecessors. It is also important to note that the triangular continuation symbol, labelled A1, is used to denote the fact that further details about the exposure of the gas line are provided in an additional sub-tree that is not shown here.

In Peterson's terms, Figure 9.9 shows how fault trees can be used to provide a general view of causality [677]. It provides some indication of the high-level failures that led to the incident. However, it is also ambiguous in the sense that there are many different reasons for the inspectors failure to report the damage to the gas line or the failure of the gas company's procedures. In contrast to Figure 9.9, Figure 9.10 provides a more singular view of the adverse occurrence. For example, it refines the abstract information in Figure 9.9 by representing the ways in which the incident developed over time. The moment at which the line coupling broke is shown to be [18:45]. Similarly, the initial UGI response is shown to have occurred during the interval between [18:50-18:58] which was too late to prevent the explosion. This representation of temporal information introduces further distinctions between our use of fault trees to support incident reconstruction and their more conventional design applications. Our approach looks at the way in which particular events actually occurred in the past rather than the probability of those events occurring again in the future. There are further complications. For example, the events in conventional fault trees tend to occur at particular instants in time. This is reflected by the way in which the developers of fault trees are encouraged to label their diagrams with 'trigger events' rather than conditions that emerge over time. For example, Andrews and Moss [27] advise that:

> "Trigger events should be coupled with 'no protective action taken', i.e. 'overheating' could be 'loss of colling' and 'no emergency shutdown'." [27]

This advise is important because it simplifies the probabilistic failure analysis that is used to guide system development. However, our application of the fault tree notation does not exploit the stochastic models that support design. As a result, it is possible to move away from any requirement for instantaneous events. For example, the Foreman's response to the initial rupture of the gas pipe took place from 18:45 to 18:54. This flexibility comes at a cost. The semantics of both the temporal information and the events in the tree become a cause for concern. For instance, Figure 9.10 uses [18:45-18:54] to denote that the Foreman coordinated a partial response to the emergency between 6.45pm and 6.54pm. In contrast, [18:58 and 19:03 approx] is used to denote the fact that two separate

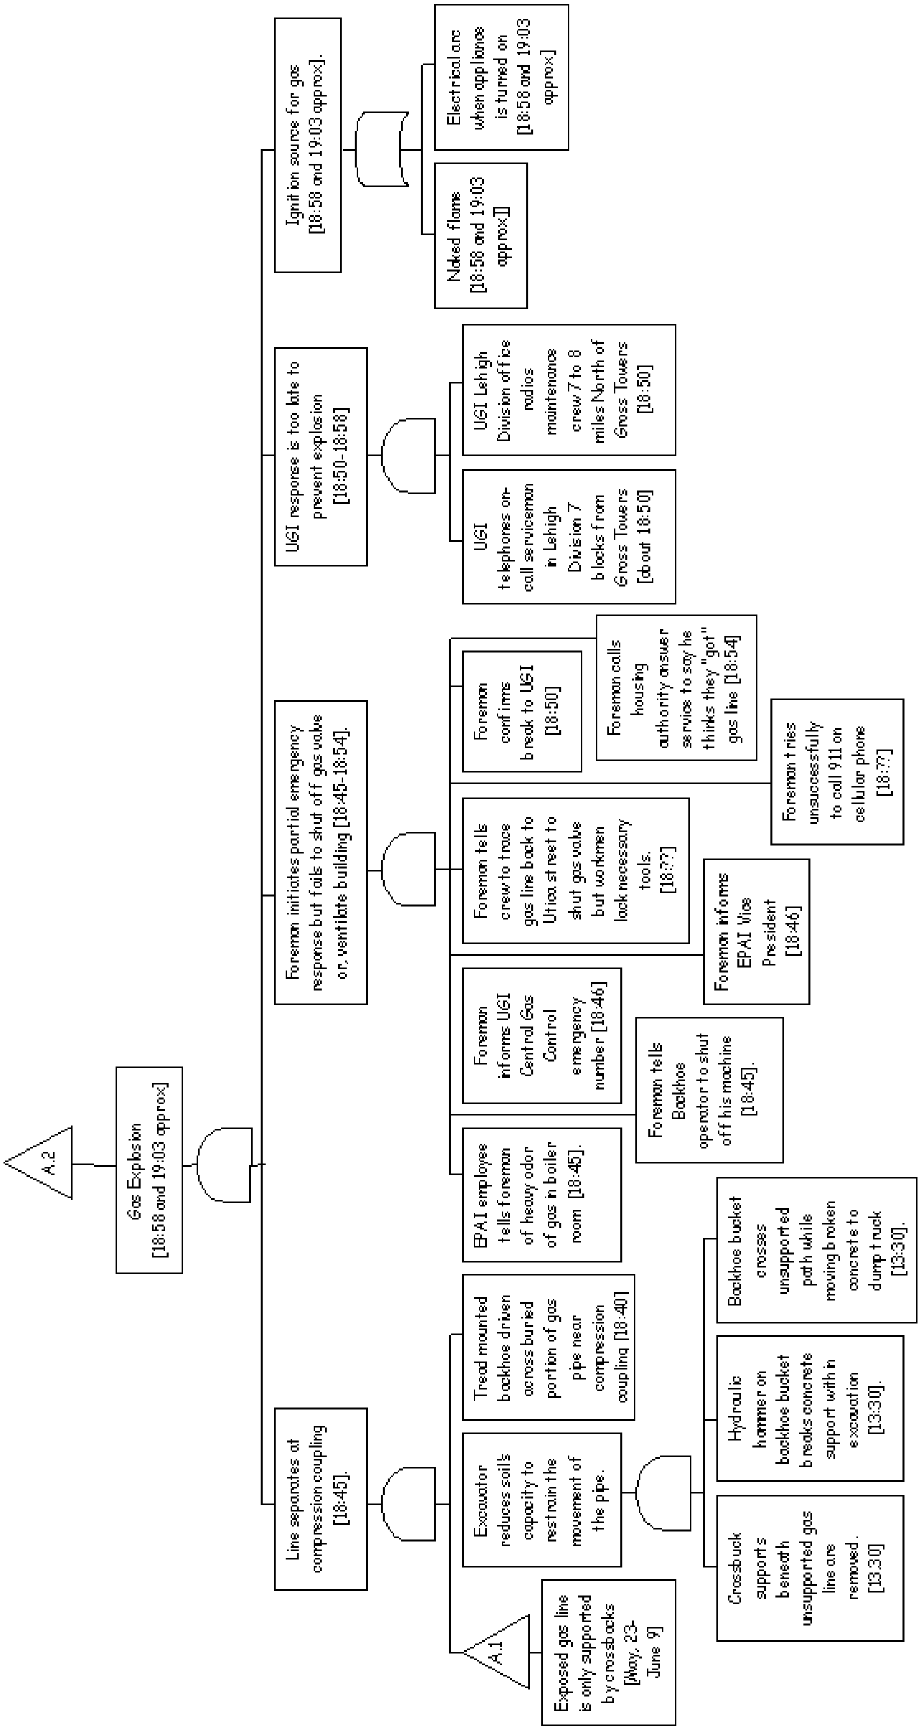
A.2
Gas Explosion [18:58 and 19:03 approx]
Line separates at compression coupling [18:45].
Foreman initiates partial emergency response but fails to shut off gas valve or, ventilate building [18:45-18:54].
UGI response is too late to prevent explosion [18:50-18:58]
Ignition source for gas [18:58 and 19:03 approx].
Exposed gas line is only supported by crossbacks [May, 23-June 9]
A.1
Excavator reduces soil's capacity to restrain the movement of the pipe.
Tread mounted backhoe driven across buried portion of gas pipe near compression coupling [18:40]
Crossbuck supports beneath unsupported gas line are removed. [13:30]
Hydraulic hammer on backhoe bucket breaks concrete support with in excavation [13:30].
Backhoe bucket crosses unsupported path while moving broken concrete to dump truck [13:30].
EPAI employee tells foreman of heavy odor of gas in boiler room [18:45].
Foreman tells Backhoe operator to shut off his machine [18:45].
Foreman informs UGI Central Gas Control emergency number [18:46]
Foreman informs EPAI Vice President [18:46]
Foreman tells crew to trace gas line back to Utica street to shut gas valve but workmen lack necessary tools. [18:??]
Foreman tries unsuccessfully to call 911 on cellular phone [18:??]
Foreman confirms break to UGI [18:50]
Foreman calls housing authority answer service to say he thinks they "got" gas line [18:54]
UGI telephones on-call serviceman in Lehigh Division 7 blocks from Gross Towers [about 18:50]
UGI Lehigh Division office radios maintenance crew 7 to 8 miles North of Gross Towers [18:50]
Naked flame [18:58 and 19:03 approx]]
Electrical arc when appliance is turned on [18:58 and 19:03 approx]

Figure 9.10: Fault Tree Showing Events Leading to Allentown Explosion

explosions occurred at 6.58pm and at approximately 7.03pm. [18:??] denotes that the exact time when the foreman attempted to call 911 is not known. These examples illustrate particular forms of temporal relationship within our case study. They are not complete in the sense that there will be temporal relationships that we cannot describe in terms of these annotations. Analysts must develop similar conventions to describe more complex timing information.

Like the graphical time-lines of the previous section, this diagram represents the passage of time flowing from left to right. For example, the lest-most sub-branch represents the events that led to the separation of the gas pipeline at 18:45. An examination of the intermediate and basic events that led to this failure shows that some, such as the initial exposure of the line, took place days before the actual failure. Other contributory events, such as the movement of the backhoe over the line occurred only minutes before the separation of the coupling. Unlike the graphical time-lines, however, this representation loosens some of the restrictions that are implied by a strict left to right ordering for events over time. It is possible to denote events that contribute to a higher level failure but for which there is little or no timing information. This is illustrated by the ambiguity that surrounds the Foreman's unsuccessful attempts to contact the emergency services by dialing '911' on his cellular telephone. No timing information is available to confirm this event because he could not raise a signal and the call was never completed.

The left to right temporal ordering of Figure 9.10 only applies to events at the same level in the tree. For instance, the basic events of the second sub-tree from the left denote that EPAI employees tell the foreman about the odour of gas and tells the Backhoe operator to stop work at 18:45. These are shown to the left of a basic event denoting the fact that the Foreman informed UGI's emergency number at 18:46 and so on. However, this left to right representation of time cannot be applied to components at different levels of the tree. For instance, an event that contributed to the separation of the gas pipeline, shown in the left-most branch, might occur *after* an event that impaired the emergency response, represented by the subtree on its right. This would, typically, occur if the inadequate response was influenced by events, such as inadequate training in emergency response procedures, that pre-dated the coupling failure.

A large proportion of the tree shown in Figure 9.10 relates to individual failures. The left-most sub-tree focuses on the excavation team's actions in exposing the gas line and in compromising the coupling. The next sub-tree deals with the Foreman's partial response to the initial separation of the gas line. However, the diagram also includes organisational factors. For example, the next sub-tree describes how UGI, the gas operating company, had only limited time to respond to the emergency. The right-most branch, in contrast, describes the environmental catalysts for two explosions. As can be seen, this sub-tree represents some of the uncertainty that inevitably arises during initial reconstructions. An inclusive OR gate shows that the explosion might have been triggered by a naked flame or by an arc from an electrical appliance.

The previous fault tree provides a graphical reconstruction of the events leading to the Allentown explosion. This offers a number of important benefits:

1. Fault trees provide an overview of the events that an analyst believes contributed to an incident. This is important because many secondary investigations gather evidence that reflects the complex nature of many safety-critical failures. It can often, therefore, be difficult to piece together evidence into a coherent account of the events that contribute to adverse incidents;

2. Fault trees also suggest alternative hypotheses and questions about the analysis that is presented in an accident report. Readers can further develop the events in a tree to develop further lines of investigation. For instance, it might be important to learn more about the problems that prevented the crew from successfully shutting off the gas flow with the tools that they had available.

Figure 9.11 introduces the events that led to the exposed gas line being supported by crossbucks. Figure 9.10 used the A1 continuation symbol to indicate the way in which these more detailed events contributed to the overall incident. In particular, it presents the more detailed events that were omitted It presents the events leading to the initial exposure of the pipeline that were denoted by the triangular extension symbol in the previous figure. The line was only supported by crossbucks

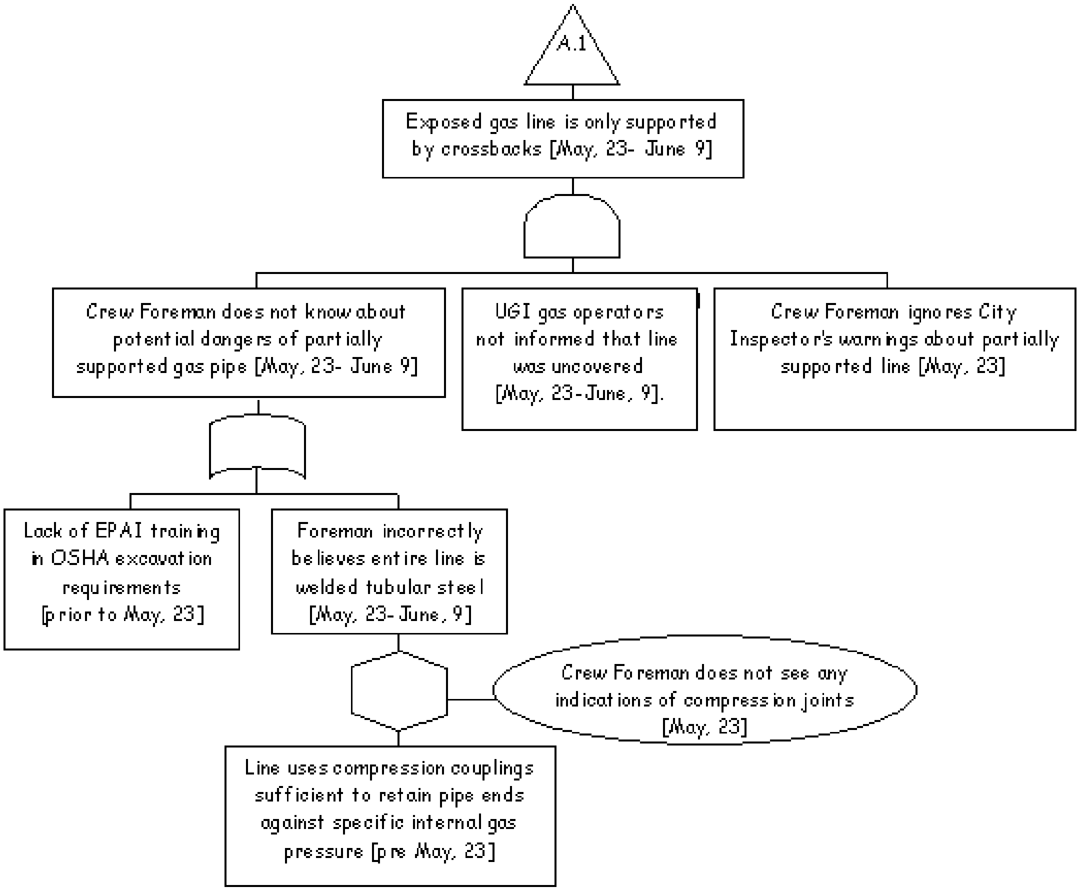
A.1
Exposed gas line is only supported by crossbacks [May, 23- June 9]
Crew Foreman does not know about potential dangers of partially supported gas pipe [May, 23- June 9]
UGI gas operators not informed that line was uncovered [May, 23-June, 9].
Crew Foreman ignores City Inspector's warnings about partially supported line [May, 23]
Lack of EPAI training in OSHA excavation requirements [prior to May, 23]
Foreman incorrectly believes entire line is welded tubular steel [May, 23- June, 9]
Crew Foreman does not see any indications of compression joints [May, 23]
Line uses compression couplings sufficient to retain pipe ends against specific internal gas pressure [pre May, 23]

Figure 9.11: Using Inhibit Gates to Represent Alternative Scenarios

because the Foreman did not appreciate the dangers of doing this, the gas supply operators, UGI, did not know that the line was uncovered and the Foreman ignored warnings from the Allentown Fire Inspectors. As can be seen, the fault tree uses an OR gate to represent a number of hypotheses about why the Foreman was unaware of the potential dangers associated with leaving the pipeline uncovered and partially supported:

> "Training–Before the accident, the workcrew had not had any formal training in excavation and trenching or in actions to take as a unit to protect lives and property in an emergency. The lack of training may account for why the crew did not shore the excavation site or tell the UGI that the gas line was unsupported. The crew foreman, despite not having any information about the construction of the gas line, said that he thought the entire line was welded tubular steel. His assumption may have led him to believe that the line could be adequately supported by crossbucks. In any event, he made a critical choice in assuming that it would be safe to leave the gas line uncovered and exposed for 2 weeks. A more prudent course of action would have been to immediately inform the UGI that the line was exposed." [588]

An OR gate is used because it is unclear what contributed most to the Foreman's lack of knowledge about the potential dangers associated with exposing the gas line. Their lack of training in appropriate OSHA standards for excavation or his incorrect belief about the pipeline construction could have affected his subsequent actions. It is important to emphasise that this decision to use an OR gate is not definitive. The construction of a fault tree is an iterative process. Subsequent discussions might discount the foreman's assumption. This might then be removed from the tree. Alternatively, it might be decided that both factors were required in order for the Foreman to behave in the way that he did. In such circumstances an AND gate might be introduced. This would have to be carefully justified because it implies that had the Foreman been trained in OSHA requirements then the incident would not have happened. Previous experience in incident and accident investigation has shown the dangers of making such assumptions about the efficacy of training as a primary protection mechanism.

In Figure 9.11 the left event of the OR gate represents the first line of analysis. It focuses on the Foreman's lack of training in applicable OSHA requirements. The second line of analysis is based on the Foreman's subsequent evidence that he believed the line to have been entirely constructed from welded tubular steel. This is developed using an INHIBIT gate, shown using a hexagon and an ellipse. The input event of an inhibit gate need not always lead to the output event. In this example, the fact that the line was constructed using compression couplings need not always lead the Foreman to incorrectly believe that an all-welded construction was used. The likelihood that the input event will lead to the output event is determined by the condition, shown in the ellipse. The Foreman did not see any indications of the compression joints and so believed that the tube was welded.

This ability to assign probabilities to representations of human error should not be underestimated. In particular, it provides a useful means of deriving simulations from a reconstruction of an incident or accident. Simulations enable analysts to replay or step through the course of an incident. Later sections will introduce automated tools for deriving simulations from incident reconstructions. For now, however, it is sufficient to observe that this can be done manually by inspecting a fault tree to trace the way in which particular combinations of events might lead to the high level failures shown in the upper levels of a figure. By introducing probabilistic information into a simulation it is possible for analysts to explore alternative scenarios during a reconstruction. For instance, Monte Carlo simulation techniques can be used to investigate probable and improbable, frequent or infrequent, traces of interaction. This approach involves the generation of random numbers typically in the range [0.0, 1.0]. This random number is then used to determine whether or not an event occurs during a particular run of the simulation. If the random number is less than the associated probability of the event then that event is assumed to happen. Conversely, if the number is greater then the event is assumed not to occur. For instance, it might be assumed that there is a 0.5 probability of anyone in the excavation team observing that compression couplings might have been used. Analysts might then begin to step through, or simulate, the events leading to the explosion. By generating a random number, it is possible to decide whether or not the couplings were observed during this particular simulation. In our example, they would be observed during approximately half of the run-throughs and would be overlooked during the rest.

Given our particular use of the fault tree notation, it might not at first appear that such simulation techniques are either appropriate or even useful. We know that the Foreman and their crew did not know that the pipe used compression couplings. However, the importance of simulation using Monte Carlo techniques is that it is possible to explore the consequences of small variations to the sequence of events that led to an incident. This is essential because incidents seldom recur in exactly the same way as previous failures. As we shall see, simulations can also be used to assess the potential impact of proposed improvements. For instance, improved training and amended site plans can be used to alert excavation crews to the construction techniques that are used by gas suppliers. These measures might increase the probability of correct observations being made to 0.8 or 0.9. It would then be correspondingly more likely that random numbers would fall below these thresholds and hence the detection of the compression joints would become more probable during any Monte Carlo simulation of a future incident. However, the obvious pitfall is that there must be some means of validating the statistics that are used to prime models such as that shown in Figure 9.11. The most appropriate means of obtaining these figures after an incident is through empirical tests with other operators. Of course, these studies are inevitably biased by the individual's knowledge that their performance is being monitored in the aftermath of an incident.

Figure 9.12 illustrates the iterative nature of incident reconstruction. This fault tree extends the diagram shown in Figure 9.11 to consider the events that contributed to the Foreman's decision

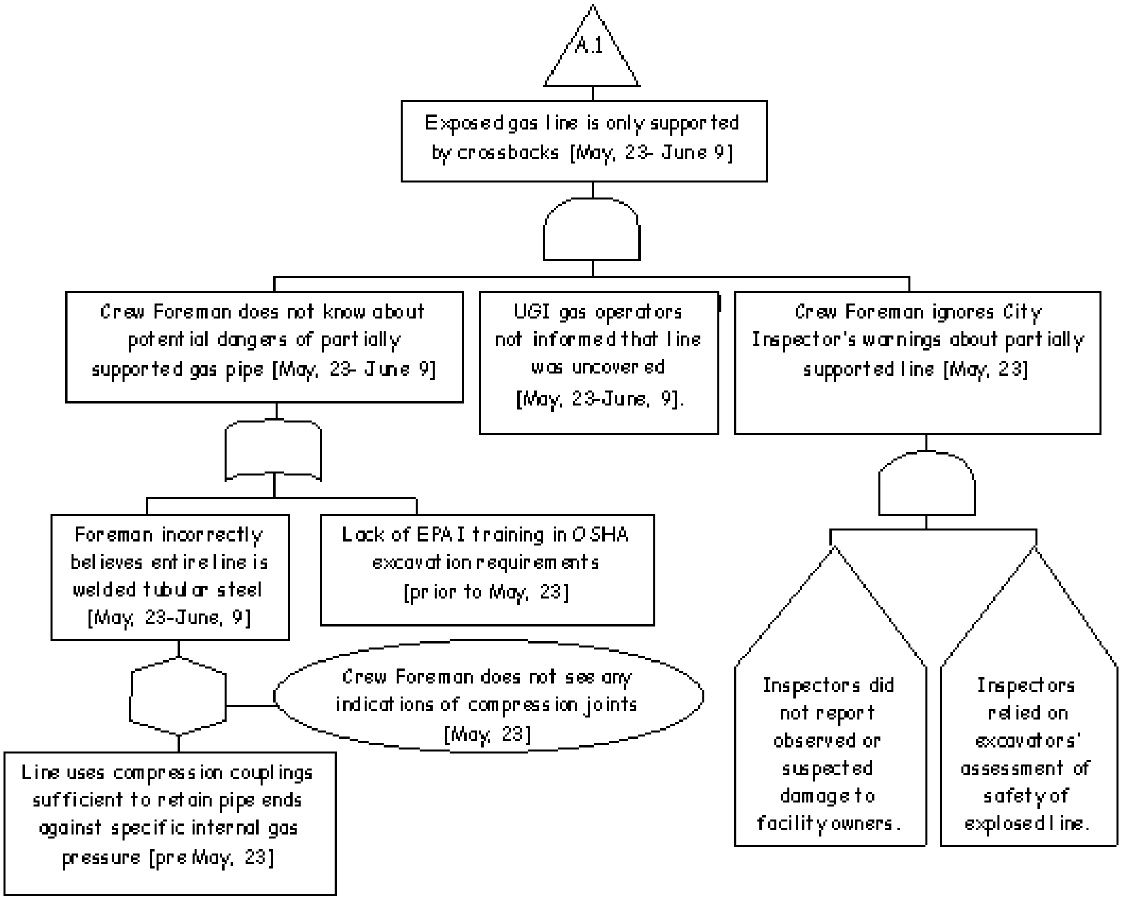
A.1
Exposed gas line is only supported by crossbacks [May, 23- June 9]
Crew Foreman does not know about potential dangers of partially supported gas pipe [May, 23- June 9]
UGI gas operators not informed that line was uncovered [May, 23-June, 9].
Crew Foreman ignores City Inspector's warnings about partially supported line [May, 23]
Foreman incorrectly believes entire line is welded tubular steel [May, 23-June, 9]
Lack of EPAI training in OSHA excavation requirements [prior to May, 23]
Crew Foreman does not see any indications of compression joints [May, 23]
Line uses compression couplings sufficient to retain pipe ends against specific internal gas pressure [pre May, 23]
Inspectors did not report observed or suspected damage to facility owners.
Inspectors relied on excavators' assessment of safety of explosed line.

Figure 9.12: Using House Events to Represent Alternative Scenarios

not to listen to the Fire Inspector's warnings. It does this by introducing HOUSE events. These are simplifications of the INHIBIT gates that were introduced in the previous chapter. HOUSE events support the simulation of alternative incident scenarios without the need to associate detailed probabilistic information with particular events. This is important because Chapter 3 has argued that it can be extremely difficult to validate human reliability statistics. In Figure 9.12, HOUSE events are used to show that the City Fire Inspectors did not report the damage to the pipeline to the facility owners and that they relied on the excavators assessment of the pipeline safety:

> "Because the citys fire inspectors saw on May 23 that the service line was unsupported, they could have prevented the accident. They showed proper concern about the safety of the line, especially after a piece of asphalt pavement fell on it and deformed it. However, not having been instructed to do otherwise, both inspectors relied on the EPAI foremen's assessment that the line was safe. It would have been more prudent of them to ask the pipeline owner for the assessment. The Safety Board concludes that the likely reason the fire inspectors did not tell the operator that its service line was damaged was because the inspectors did not understand the importance of notifying operators so the effects on a facility could be assessed by the operators and necessary action taken. Had the inspectors notified the UGI, it, the Safety Board believes, would have taken the necessary corrective actions, and the accident would not have happened." [588]

HOUSE events can either be "turned" on or off during the analysis of a fault tree. The NTSB investigation found that the Inspectors failed to report the damage and that they relied on the excavators. Technically, this can be represented by assigning a probability of 1 to the two house events in Figure9.12. However, the ability to switch events on and off also provides analysts with means of exploring alternative hypotheses about the course of an accident. For instance, a house event can be turned off if it is assigned a probability of 0. This can be used to explore what might have happened if the Inspector had reported the damage to the pipeline or had performed their own assessment of the pipeline safety. This might then have prevented the Foreman from ignoring their initial warnings about the unsupported line.

The previous paragraphs have argued that fault trees can be used to provide an overview of the immediate human errors that contribute to incidents. House events and inhibit gates can also be used to analyse the factors that did not play a part in past failures but which might lead to similar errors during the future operation of the system. In contrast, Figure 9.13 extends the previous analysis to look beyond the explosion at the emergency response. The continuation symbol, A2, is used to indicate that the events leading to the explosion, shown in Figure 9.10, also form part of this tree. In contrast, however, Figure 9.13 illustrates the events that contributed to an effective and well-co-ordinated response. This is an important illustration of how a graphical notation can provide a high-level overview of both the failures that contribute to an incident and the mitigating factors that help to reduce its potential consequences. Some of these events stem from successful training and management:

> "The fire department used the city's mass casualty incident plan, and the coordinator used the fire department's incident command system. The command post was established on the front lawn of Gross Towers at 7:03; and at 7:04, the emergency-response staging area and emergency shelter were established at the Allentown Fairgrounds, about 1/2 mile southwest of Gross Towers, where approximately 200 residents and 150 family members were helped. At 7:21, a MedEvac helicopter was requested to transport burn victims. Buses were requested at 7:40 to transport victims to the shelter at the fairgrounds, and by 7:49, the preliminary search of Gross Towers for victims was complete. The last injured resident was transported to a local hospital at 8:45." [588]

Other events that contributed to an effective and well-coordinated response were more due to chance than to planning. For example, the fact that many residents were not in the building at the time of the explosion helped to reduce the demands on those coordinating the initial evacuation. As can be seen from Figure 9.13, these 'chance' factors are not explored to the same level of detail as the organisational successes. This, in part, reflects the amount that can be gained from an improved understanding of these different aspects of the incident. It could also be argued that such 'chance' events ought to be denoted by HOUSE events so that analysts do not assume that they will always be true during any subsequent simulations of similar incidents.

Moving from Reconstructions to Conclusions

The previous fault trees, with the exception of Figure 9.9, illustrate the way in which the graphical notation can reconstruct the events leading to an incident. Fault trees provide a mid point between the evidence from any secondary investigation and the causal analysis that is the focus of the next chapter. The difference between reconstruction and causal analysis is often embodied in the structure of incident reports. For example, the NTSB report into the Allentown incident contained separate sections entitled 'Investigation', which includes the reconstruction of the events leading to the incident, and 'Analysis', which uses the reconstruction to support arguments about the underlying causes of the explosion. The findings of the analysis help to shape the conclusions that are to be drawn from any investigation. The following quotations illustrate these differences:

> "It took about 6 hours for the hydraulic hammer to break the concrete up. According to the EPAI employees, the impact of the hammer caused the ground to vibrate significantly. The backhoe bucket was used to remove the broken concrete and to load the pieces into a dump truck. The path of the backhoe bucket crossed over the pipe. The backhoe operator said that about 6:40 p.m. he moved the backhoe from a spot south of the

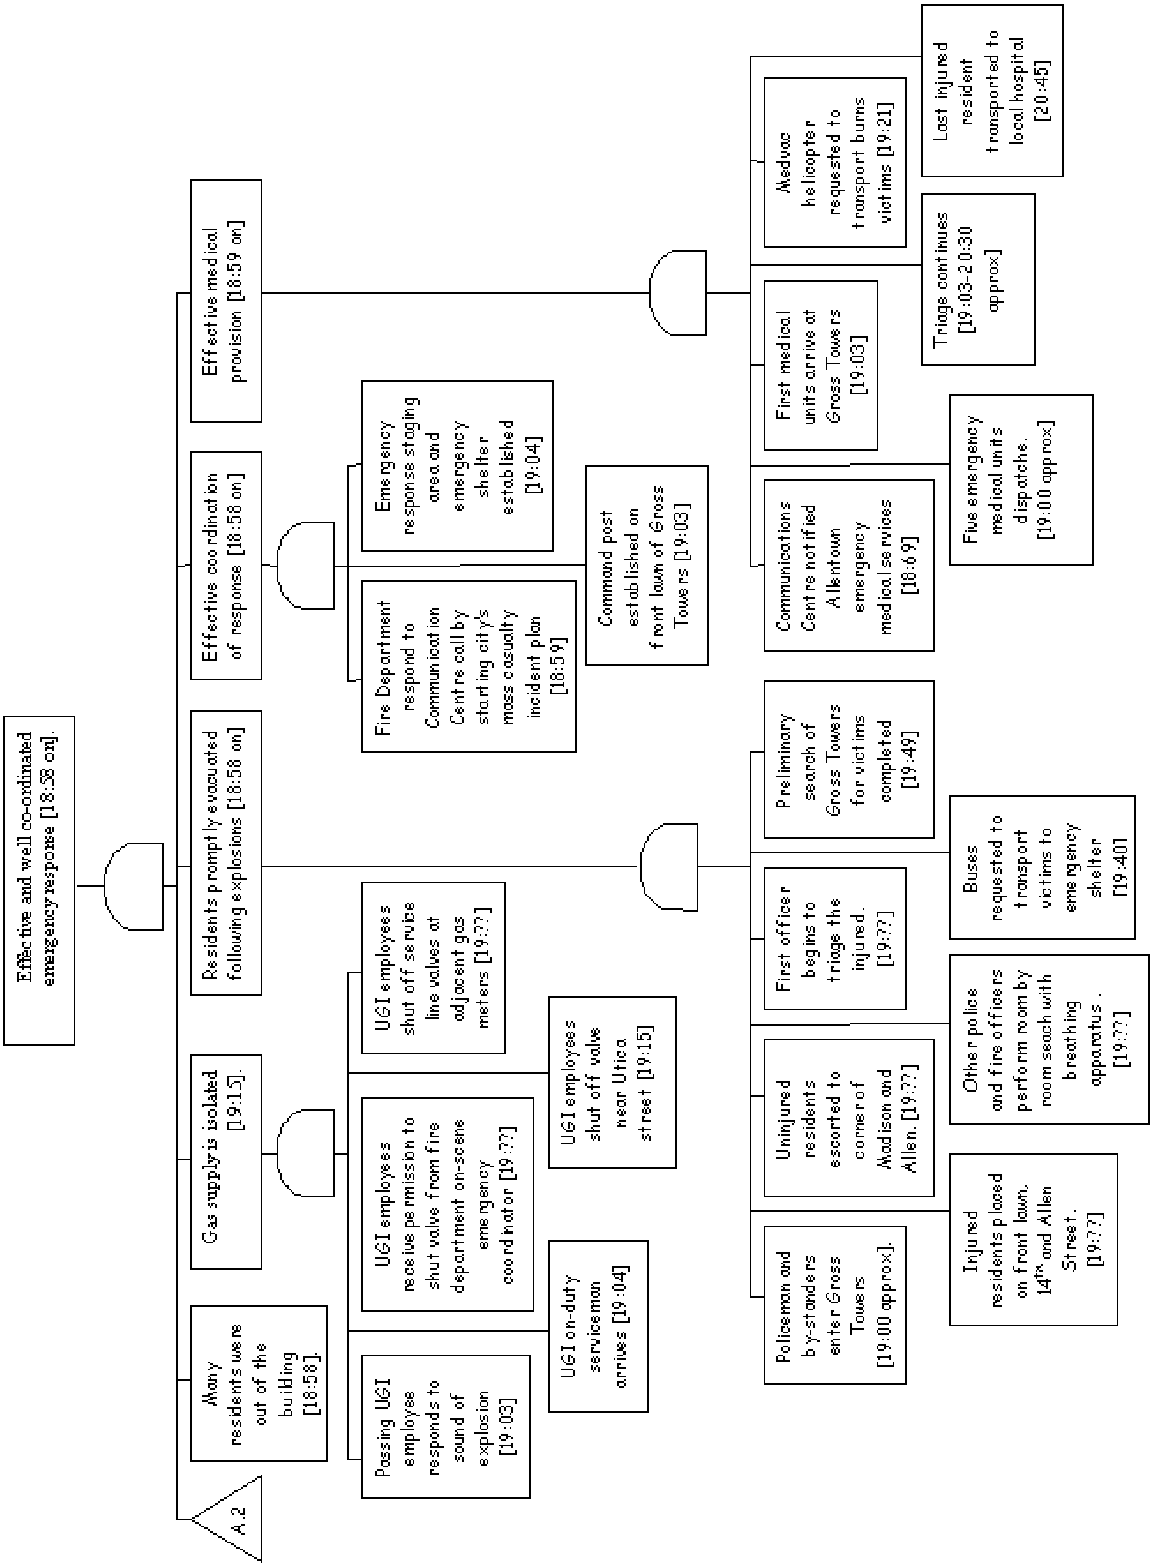
Effective and well co-ordinated emergency response [18:58 on].
A.2
Many residents were out of the building [18:58].
Gas supply is isolated [19:15].
Residents promptly evacuated following explosions [18:58 on]
Effective coordination of response [18:58 on]
Effective medical provision [18:59 on]
Passing UGI employee responds to sound of explosion [19:03]
UGI on-duty serviceman arrives [19:04]
UGI employees receive permission to shut valve from fire department on-scene emergency coordinator [19:??]
UGI employees shut off valve near Utica street [19:15]
UGI employees shut off service line valves at adjacent gas meters [19:??]
Fire Department respond to Communication Centre call by starting city's mass casualty incident plan [18:59]
Command post established on front lawn of Gross Towers [19:03]
Emergency response staging area and emergency shelter established [19:04]
Policeman and by-standers enter Gross Towers [19:00 approx].
Injured residents placed on front lawn, 14th and Allen Street. [19:??]
Uninjured residents escorted to corner of Madison and Allen. [19:??]
Other police and fire officers perform room by room seach with breathing apparatus . [19:??]
First officer begins to triage the injured. [19:??]
Buses requested to transport victims to emergency shelter [19:40]
Preliminary search of Gross Towers for victims completed [19:49]
Communications Centre notified Allentown emergency medical services [18:69]
Five emergency medical units dispatche. [19:00 approx]
First medical units arrive at Gross Towers [19:03]
Triage continues [19:03-20:30 approx]
Medvac helicopter requested to transport burns victims [19:21]
Last injured resident transported to local hospital [20:45]

Figure 9.13: Fault Tree Showing Post-Explosion Events

excavation to one on the west. In moving it, he crossed a buried section of pipeline that was between the excavation and the north wall of Gross Towers. The odour of gas was first detected about 6:45 p.m."
([588], Investigation, page 11).

"When the excavator resumed on June 9, its activities near the service line probably reduced the amount of restraint provided by the soil even more and increased the longitudinal force enough to cause the pipe to separate fully from the coupling. Using the impact tool to break the concrete tank support and moving the backhoe over the pipeline caused the soil to vibrate and probably further reduced the soils restriction of pipe movement. Also, the backhoe probably struck the line when being operated across it; the foreman's reports to both the UGI and the housing authority indicated that the pipe had been struck during recent excavation activities. Although the foreman denied after the accident that the backhoe had struck the line, the coating of the pipe showed evidence of mechanical damage, as did the pipe steel at one location. Also, the foreman's calls both to the housing authority and to the UGI show that at the time he believed his crew had hit the gas line while excavating."
([588], Analysis, page 32).

"By reducing the soils capacity to restrain the movement of the pipe and by exerting forces on the service line that resulted in excessive longitudinal stress, the excavator caused the line to separate at a compression coupling."
([588], Conclusion, page 47).

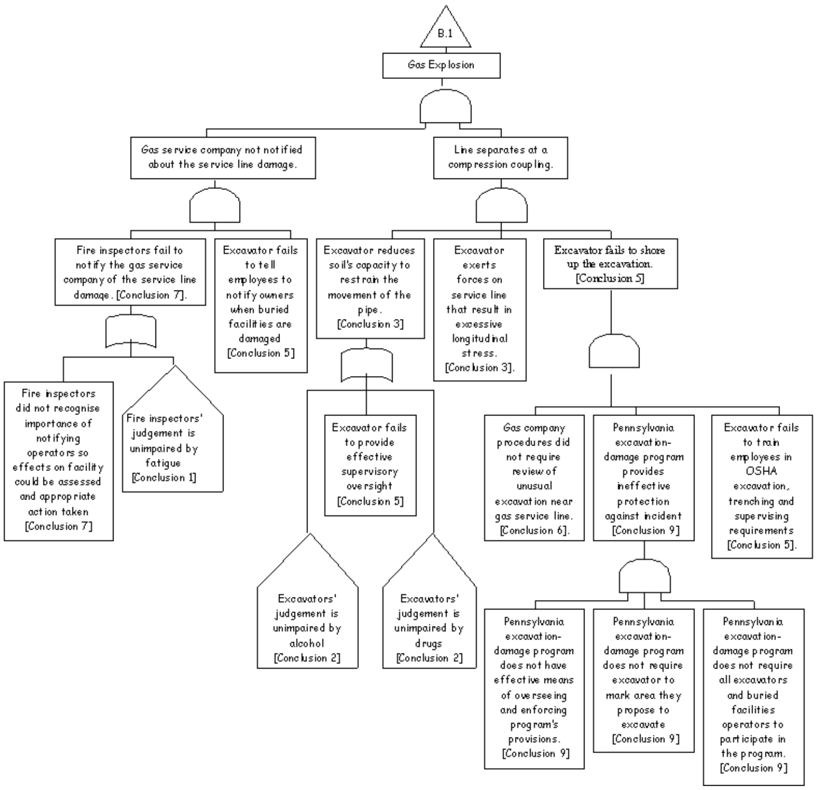
B.1
Gas Explosion
Gas service company not notified about the service line damage.
Line separates at a compression coupling.
Fire inspectors fail to notify the gas service company of the service line damage. [Conclusion 7].
Excavator fails to tell employees to notify owners when buried facilities are damaged [Conclusion 5]
Excavator reduces soil's capacity to restrain the movement of the pipe. [Conclusion 3]
Excavator exerts forces on service line that result in excessive longitudinal stress. [Conclusion 3].
Excavator fails to shore up the excavation. [Conclusion 5]
Fire inspectors did not recognise importance of notifying operators so effects on facility could be assessed and appropriate action taken [Conclusion 7]
Fire inspectors' judgement is unimpaired by fatigue [Conclusion 1]
Excavator fails to provide effective supervisory oversight [Conclusion 5]
Gas company procedures did not require review of unusual excavation near gas service line. [Conclusion 6].
Pennsylvania excavation-damage program provides ineffective protection against incident [Conclusion 9]
Excavator fails to train employees in OSHA excavation, trenching and supervising requirements [Conclusion 5].
Excavators' judgement is unimpaired by alcohol [Conclusion 2]
Excavators' judgement is unimpaired by drugs [Conclusion 2]
Pennsylvania excavation-damage program does not have effective means of overseeing and enforcing program's provisions. [Conclusion 9]
Pennsylvania excavation-damage program does not require excavator to mark area they propose to excavate [Conclusion 9]
Pennsylvania excavation-damage program does not require all excavators and buried facilities operators to participate in the program. [Conclusion 9]

Figure 9.14: Fault Tree Showing NTSB Conclusions about the Causes of the Explosion

The structure of the NTSB report separates the presentation of reconstruction, causal analysis and conclusions. We have, however, argued that these different activities often become blurred during the process of incident investigation. The reconstruction of an incident inevitably involves the formation and testing of causal hypotheses. Investigators include events in a reconstruction because they believe that those events have had some impact on the course of an incident. For example, if it were believed that the timing of the foreman's 911 call was critical for the analysis of the Allentown explosion then evidence would be sought so that this event could be explicitly included in any reconstruction. If the attempted call was not thought to have a significant, or potential, impact then it might be omitted. The generation and testing of such causal hypotheses against any reconstruction will inevitably affect the conclusions that can be drawn from an investigation. These links make it important that any tools, including time-lines and fault trees, do not impair the complementary activities of reconstructing an incident, generating causal hypotheses and forming conclusions.

Figure 9.14 shows how fault trees can be used to summarise the conclusions from the NTSB's investigation into the Allentown incident. Such high-level overviews are important because they help to determine whether the individual findings of an investigation form a coherent argument. For example, Figure 9.14 shows how the excavators' failure to shore-up the excavation was not simply due to individual failure on the part of the foreman and his team. The NTSB investigators also identified higher-level failures on the part of the gas company, on the excavations company and on the Pennsylvania excavation-damage program. Figure 9.14 shows how fault trees can be used to explicitly represent the relationships between these individual conclusions. The NTSB's organisational and managerial conclusions in Figure 9.14 contrast with OSHA's findings about the health and safety aspects of this incident. OSHA focuses more narrowly on the individual human errors that were represented in previous reconstructions, such as Figure 9.10:

> "OSHA determined that the EPAI foreman did not meet OSHAs definition of competence, as stated in 26 CFR 1926.650 (b). Among the failures OSHA attributed to the foreman were that he had classified the soil type incorrectly, had improperly supported the gas line, did not recognize the hazard of the gas line, did not know the lifting capacity of the chain used in the failed attempt to lift the fuel tank, did not know the lifting capacity of the backhoe, and did not keep spoil from the excavation from the top edge of the excavation." [588]

Before proposing further benefits that can be derived from using fault trees to reconstruct and summarise the conclusions of an incident investigation, it is important to acknowledge a number of weaknesses. Previous sections have argued that these is no automatic means of moving from the evidence of primary and secondary investigations to the reconstructions of Figures 9.10 and 9.11. Similarly, there is no automatic means of moving from incident reconstructions, such as Figures 9.10 and 9.11, to the conclusion overview presented in Figure 9.14. Both activities rely upon the skill and experience of individual analysts. Fault trees are, therefore, not a panacea. They simply provide a means of representing and reasoning about the products of different stages in an incident investigation.

The lack of any automated means of moving between fault tree reconstructions, illustrated in Figure 9.13, and conclusions, illustrated by Figure 9.14, should not be surprising. As we have seen, reconstructions tend to focus on the proximal events surrounding a particular incident. For example, Figure 9.10 traces the way in which initial failures on the 23rd May led to the eventual explosion in Allentown on June, 9th. However, many incident reports combine findings about specific causes with conclusions about wider failures in the managerial and regulatory system. For instance, Figure 9.14 considers problems at a State level, through the failure of the excavation damage program, and at a national level, through the lack of OSHA training for excavation workers. Hence the conclusions of an incident report are likely to draw on information that is not, typically, included within the reconstruction of a single incident.

There are further, more theoretical barriers to the automatic generation of conclusions from reconstructions. Previous chapters have argued that the interpretation and analysis of evidence is influenced by the goals and priorities of the organisations that are involved in an investigation.

Most often this is interpreted as a 'bad thing'. Organisations seek to influence or bias the findings of an investigation for commercial and even political ends. However, the social processes of incident investigation can also have a positive effect. For instance, regulators often increase the salience of particular pieces of evidence if they support the findings of previous incident reports. This is illustrated by the NTSB's emphasis on the importance of excess flow valves following the Allentown explosion. This was seen to be yet another example of an incident that might have been mitigated by the use of these devices. As a result the conclusions of the report places the Allentown incident in the context of many previous incidents that could not be explicitly considered within a reconstruction of this particular incident:

> "In the past 20 years, the Research and Special Programs Administration has failed to effectively assess the benefits of excess flow valves and has failed to promote their use." ([588], Conclusions, page 48).

Any system that attempted to generate conclusions from a reconstruction would also have to consider the wider commercial, political and regulatory environment in which it was operating. Although incident investigators must be independent from industry regulators, it is important that they work together to push through the recommendations of any enquiry. Ultimately, regulators are free to reject the findings of an investigation if they do not believe that they would lead to safety improvements. This need for independence and cooperation poses considerable social, organisational and technical challenges.

A more serious criticism of the fault tree notation, illustrated in Figure 9.14, is that it fails to distinguish between contextual and contributory factors and the root causes that were introduced in Chapter 7. Andrews and Moss maintain that fault trees are intended to record the "immediate, necessary and sufficient" events that contribute to any failure [27]. As a result, almost every conclusions represented in Figure 9.14 is elevated to the status of a root cause. There is no way of representing the observation that Pennsylvania's ineffective excavation damage program might have *contributed* to the incident but did not directly *cause* it. Such distinctions might be represented by introducing additional syntactic features into the basic fault tree notation. However, this would sacrifice many of the benefits associated with the use of an existing and well understood notation. Chapter 10 will explore these issues in greater detail. For now it is sufficient to observe that although it is possible to use fault trees to provide an overview both of the events leading to an incident and of the conclusions that can be drawn from an incident, there remain a number of theoretical and practical barriers to this application of the existing notation.

Figure 9.14 focussed on the failures that led from the damaged pipeline to the eventual explosion. In contrast, Figure 9.15 shows how the consequences of this incident were largely determined by the response *after* the pipeline was damaged. For example, the Allentown investigation found that the city's mounted an effective response to this incident. Careful preparation and training were guided by the lessons of previous incidents:

> "The executive director stated that the housing authority had procedures for evacuating the occupants and that the residents practiced the routines. For example, every 6 months the fire department conducted fire inspections and drills that also tested the evacuation procedures and emphasized how important it was for the residents to respond promptly. The drills included special precautions for the elderly and handicapped; and after a drill was held, all residents participated in a critique. Placards were posted on the windows and doors of apartments that had handicapped occupants and of rooms in which occupants were using pressurised oxygen." [588]

Figure 9.15 uses a HOUSE event to represent the finding that the housing association and the city's emergency response were appropriate. Previous sections have shown how these events can be 'turned' on or off during any walk-through of the causal model. As a result, analysts are encouraged to hypothesise about the potential impact of an ineffective response. In Figure 9.15, this would indicate a failure to learn from previous incidents and ultimately would have contributed to injury and a loss of life during the incident.

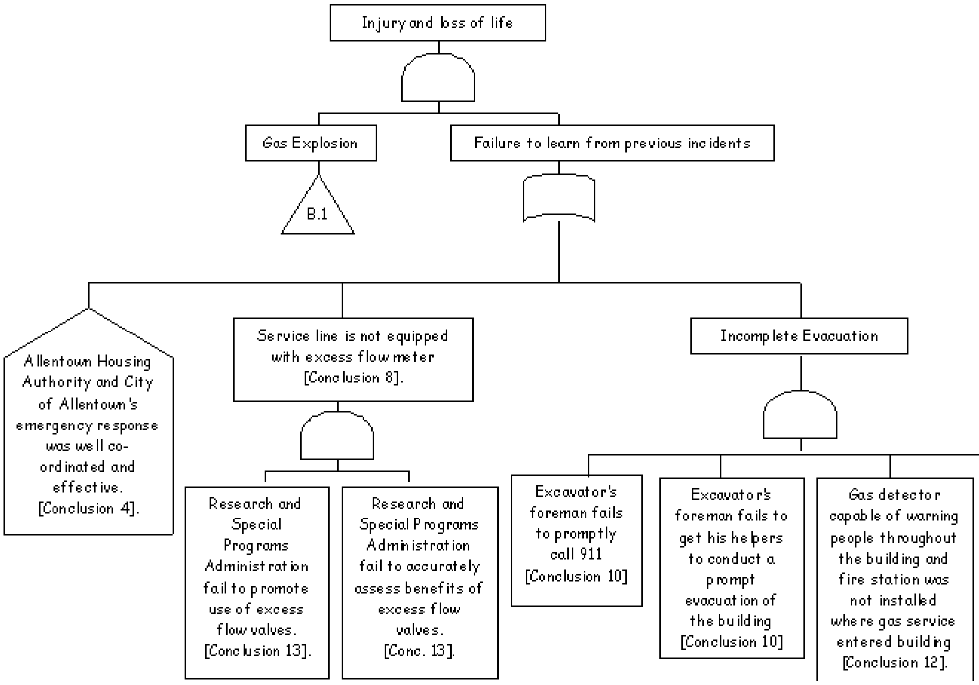
Injury and loss of life
Gas Explosion
B.1
Failure to learn from previous incidents
Allentown Housing Authority and City of Allentown's emergency response was well co-ordinated and effective. [Conclusion 4].
Service line is not equipped with excess flow meter [Conclusion 8].
Research and Special Programs Administration fail to promote use of excess flow valves. [Conclusion 13].
Research and Special Programs Administration fail to accurately assess benefits of excess flow valves. [Conc. 13].
Incomplete Evacuation
Excavator's foreman fails to promptly call 911 [Conclusion 10]
Excavator's foreman fails to get his helpers to conduct a prompt evacuation of the building [Conclusion 10]
Gas detector capable of warning people throughout the building and fire station was not installed where gas service entered building [Conclusion 12].
Excavation crew members are not trained in handling an emergency [Conclusion 11]

Figure 9.15: Fault Tree Showing Conclusions about Injuries and Loss of Life

The previous analysis raises many questions about the role of organisational failure in incidents and accidents. For instance, Figure 9.15 suggests that the lack of an excess flow valve or meter is an indication of a failure in organisational learning. As we have seen, the NTSB investigators argued that this stemmed from the Research and Special Programs Administration's failure to promote or to accurately assess the benefits of these devices. However, it is not certain that such devices will always prevent incidents such as the Allentown explosion. This objection can be represented by replacing the basic events in Figure 9.15 with an inhibit gate, as shown in Figure 9.16. Analysts could then assign a probability to the likelihood that an EFV would have cut the supply of gas either before or after the explosion. It might seem that it would be a trivial exercise to derive such reliability data given modern testing methods. Certainly, it ought to be easier to assess the reliability of such devices than it is to quantify human reliability assessments. As we have seen, however, the economic consequences of requiring the introduction of EFVs led to considerable debate about their reliability and utility between the supply industry and their regulators:

> "The two-accident sample RSPA (Research and Special Programs Administration within the Department of Transportation responsible for pipeline safety) used in its 1995 study to assess EFV effectiveness is statistically insignificant. Even so, RSPA incorrectly assessed what happened in the two accidents it did use. Although a life was saved when an EFV operated properly in one of the accidents, RSPA attributed its benefit as only one fifth of the $ 2.6 million used by the study as the value of a life. That error was further compounded by using 57 percent as an assumed EFV effectiveness percentage. When Safety Board representatives met with RSPA on March 16, 1995, it questioned RSPA about the basis for the effectiveness percentage. A RSPA economist explained that 95 percent effectiveness was initially used, but that number was reduced because a National Highway Traffic Safety Administration (NHTSA) analyst, not knowledgeable about EFVs, said he believed the number was to high. RSPA stated that even though it had no justification for a different percentage, it offered 57 percent as the effectiveness percentage, and the NHTSA analyst accepted it, saying that it seemed about right.

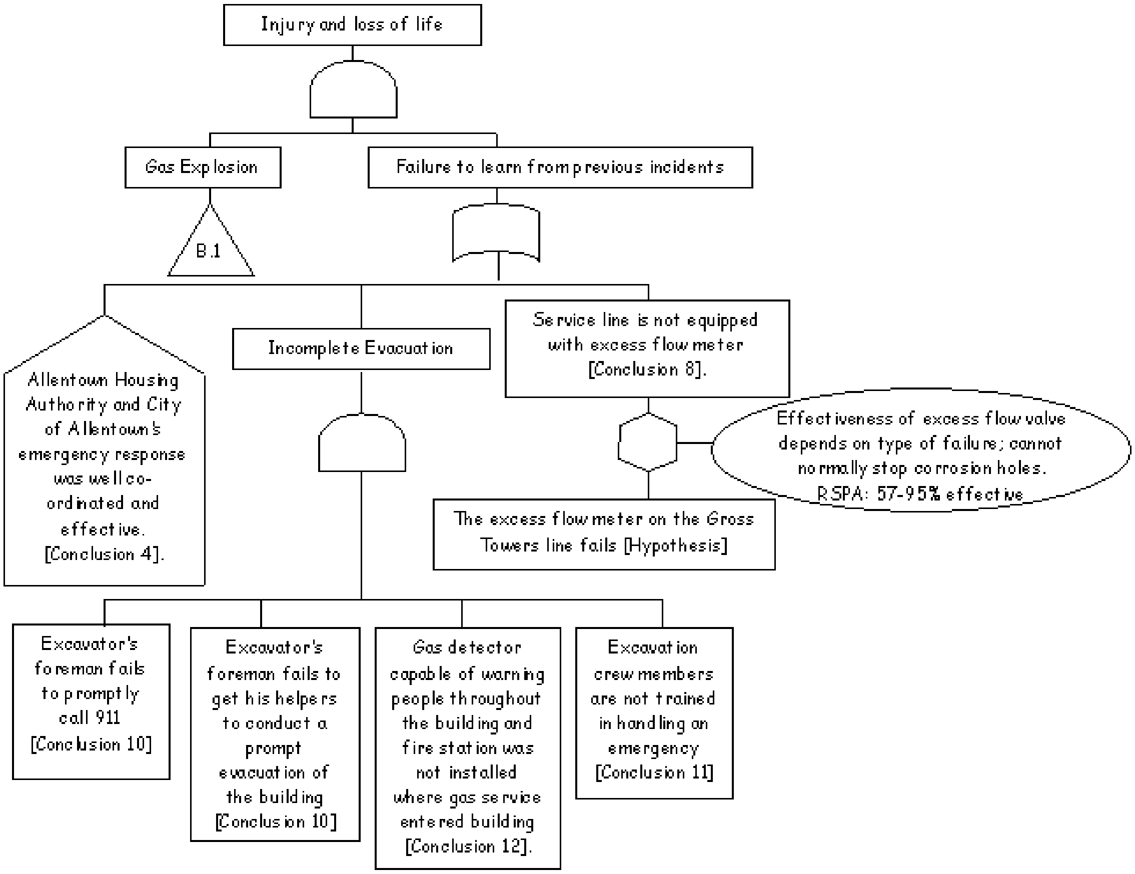
Injury and loss of life
Gas Explosion
Failure to learn from previous incidents
B.1
Allentown Housing Authority and City of Allentown's emergency response was well co-ordinated and effective. [Conclusion 4].
Incomplete Evacuation
Service line is not equipped with excess flow meter [Conclusion 8].
Effectiveness of excess flow valve depends on type of failure; cannot normally stop corrosion holes. RSPA: 57-95% effective
The excess flow meter on the Gross Towers line fails [Hypothesis]
Excavator's foreman fails to promptly call 911 [Conclusion 10]
Excavator's foreman fails to get his helpers to conduct a prompt evacuation of the building [Conclusion 10]
Gas detector capable of warning people throughout the building and fire station was not installed where gas service entered building [Conclusion 12].
Excavation crew members are not trained in handling an emergency [Conclusion 11]

Figure 9.16: Fault Tree Showing Conclusions about Reliability of Excess Flow Valves

> Other parts of RSPA's study appear to include similar insupportable numbers and assumptions." [588]

This quotation illustrates the way in which reliability data assumes a particular social and organisational significance in the aftermath of an incident. It is important to emphasise that quantitative reliability assessments are not always objective and that their true value is often questioned in the aftermath of an adverse occurrence.

9.1.3 Petri Nets

The previous section has shown how Fault-trees can be used to reconstruct the events that lead to incidents and accidents. We have also shown how they can be used to provide an overview of the conclusions that emerge from the subsequent analysis of those reconstructions. However, we have also noted a number of limitations in using this notation to distinguish between root causes and other contributory or contextual factors. The European Federation of Chemical Engineering's International Study Group On Risk Analysis also concludes:

> "Fault-trees have difficulties with event sequences... parts of systems where sequence is important are, therefore, usually modelled using techniques more adept at incorporating such considerations" [188].

We have tried to address this criticism by annotating events with real-time labels. However, this creates additional problems for analysts who must represent the way in which many failures emerge over a prolonged period of time. For example, the Allentown pipeline was left with inadequate support from the 23rd May until the 9th June. The following pages, therefore, introduce an alternative graphical notation that can be used to reconstruct the events that contribute to safety-critical incidents.

Petri Nets were developed to support the engineering of concurrent systems [460]. Chretienne shows how they can be used to represent and reason about timing properties of different systems designs [165]. Some notable attempts have been made to represent human factors requirements using this notation. For instance, Van Biljon exploits Petri Nets to derive formal specifications of interactive systems at a very high level of abstraction [81]. Bastide and Palanque have used this notation to represent the design of an interactive database [69, 663]. Hura and Attwood have used Petri Nets to support accident analysis from the perspective of hardware and software engineering engineering [377]. In contrast, this sections uses the same notation to reconstruct the more general systems failures that characterise safety-critical incidents.

A number of limitations complicate the application of Petri Nets to analyse accidents that involve interactive systems. In particular, they do not capture 'real' time. Various modifications have been applied to the classic model. Levi and Agrawala use 'time augmented' Petri Nets to introduce the concept of 'proving safety in the presence of time' [488]. Unfortunately, these enhancements are too complex to provide practicable tools for incident analysis. The following pages, therefore, retain a most basic form of the Petri Nets notation. It should be noted, however, that a range of modelling tools are significantly reducing the burdens associated with more advanced, time-augmented and stochastic extensions.

Petri Nets have been specifically developed to represent the complex sequencing and synchronisation constraints that cannot easily be captured by fault trees and time-lines. They can be used to reconstruct an incident in terms of the conditions that are satisfied at particular moments [678]. These conditions together help to represent the state of the various systems, individuals and groups that are involved in an adverse occurrence. The state of these diverse and distributed components will change during the course of an incident. Petri Nets model this by representing the way in which certain events can occur if particular conditions hold. If an event takes place then it can alter the state of the people, systems etc involved in the incident. Changes in state are represented by the new conditions that hold after an event has occurred. These new conditions enable further events to take place.

Places can be used to describe the conditions which hold for operators and their systems during the course of an incident. In our case study, investigators might use a place to represent the fact that the gas line is exposed. Another place can represent the fact that the excavation is undertaken on the incorrect assumption that the soil has a compression strength of 1.5 tons per square foot. Such places describe the causes of an incident at an extremely high level of abstraction. Places can also be used to represent causes which are specifically related to the human factors or systems engineering of an application. Places can be used to represent human factors observations about the behaviour of individual operators; the Allentown Fire Inspector is concerned about the consequences of the land slide. They can represent environmental attributes, such as the soil around the tank that is being extracted is contaminated with fuel. Places might also represent the behaviour of individual systems; the hydraulic hammer is breaking up the concrete base.

Transitions can be used to represent the events that trigger incidents and accidents. The initiating event leading to the Allentown explosion can be identified as the Foreman's over-estimate about the potential strength of the soil that he was excavating:

> "The foreman evaluated the soil being excavated as OSHA Type A, which is cohesive soil with an unconfined compressive strength of 1.5 tons per square foot. (OSHA's post-accident evaluation indicated that a visual evaluation of the soil should have shown that it was OSHA Type C, which is a cohesive soil with an unconfined compressive strength of 0.5 ton or less per square foot.) While an Allentown inspector was inspecting the EPAI's work, he saw the excavation's west sidewall slide into the excavation exposing the gas line, which was about 3 to 4 feet west of the tank. The collapsed sidewall removed the soil support from about 30 feet of the gas line, causing it to sag." [588]

The foreman's over-estimate of the soil strength can be represented as a transition that changes the state of the wider 'system' into one in which an excavation proceeds with inadequate precautions. This can be represented as a place that, if marked, can lead to a further transition, which triggers the land slip. Isolating these critical transitions provides a focus for subsequent analysis. In particular, the previous analysis might provoke greater discussion of the reasons why the foreman made an incorrect assessment of the soil strength.

Petri Nets have a formal syntax and semantics. The structure of valid networks and the meaning of those networks can be precisely defined using relatively simple mathematical concepts. Petri Nets are directed graphs; $PN = (P, T, E, M)$. They consist of a set of places, P, transitions, T, edges, E and markings, M. Edges connect places to transitions: $E \subseteq \{P \times T\} \cup \{T \times P\}$. They can be used to form the chains of events and conditions that lead to an accident. They can be described in terms of two functions. The function *Op* maps from each transition to its set of output places. The output places of a transition represent the conditions which hold after an event has occurred. For example, an output place can be used to represent the observation that the gas line is exposed after the land slip has occurred. An input place function, *Ip*, maps from each transition to the set of input places for that transition. The input places of a transitions specify the conditions which must hold for an event to occur. The input place of a transition can be used to represent the observation that the incorrect assumption about soil strength during the excavation led to the soil slip.

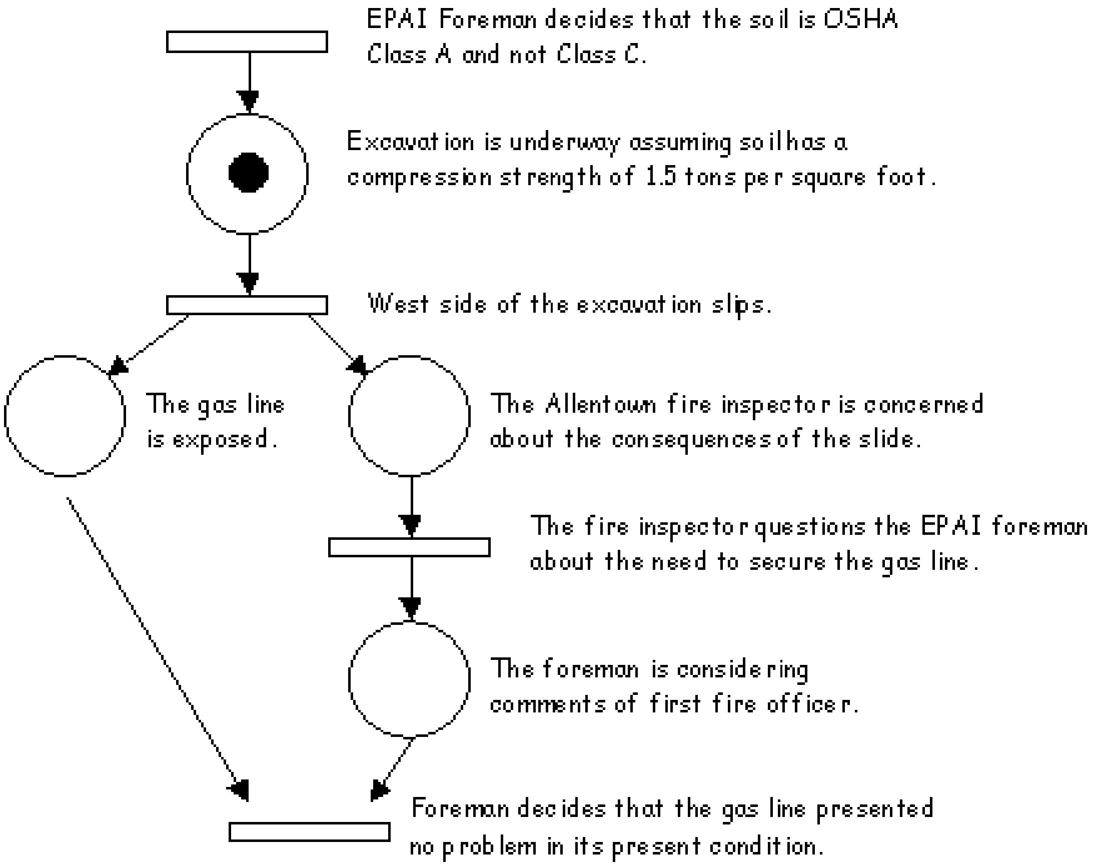
EPAI Foreman decides that the soil is OSHA Class A and not Class C.
Excavation is underway assuming soil has a compression strength of 1.5 tons per square foot.
West side of the excavation slips.
The gas line is exposed.
The Allentown fire inspector is concerned about the consequences of the slide.
The fire inspector questions the EPAI foreman about the need to secure the gas line.
The foreman is considering comments of first fire officer.
Foreman decides that the gas line presented no problem in its present condition.

Figure 9.17: Petri Net of Initial Events in the Allentown Incident

Fortunately for those who are more interested in the application than in the formal underpinnings of this notation, Petri Nets also have a graphical representation. Events, or transitions, are shown as bars (–). Conditions, or places, are denoted by unfilled circles (○). Edges are shown as arrows linking places and transitions. Figure 9.17 shows how a Petri net can represent the events leading up to the Allentown incident. The filled in circles represent tokens. These 'mark' the unfilled circles, or places, that represent assertions about th e state of the system. In this diagram, a place is marked to show that the excavation is underway assuming that the soil has a compression strength of 1.5 tons per square foot. An important benefit of the Petri Net notation is that analysts can simulate the flow of events in an accident model by altering the markings in a network. This is done through an iterative process of marking and firing. If all of the places leading to a transition, denoted by the rectangles, are marked then that transition can fire. In Figure 9.17, the transitions labelled West side of excavation slips can fire. All of the output places from this transition will then be marked. For example, if the place labelled West side of excavation slips were to fire then the places The gas line is exposed and The Allentown fire inspector is concerned about the consequences of the slide would be marked and the tokens in places that triggered this transition would be removed.

In order to simulate the dynamic events during an incident, tokens are used to mark those places in a Petri Net which are enabled. A place is enabled if its conditions hold. The tokens in a net are said to characterise a marking state and are denoted graphically by filled dots (•). For instance, Figure 9.18 is marked to show that the gas line is exposed and that both of the Allentown Fire Inspectors are concerned about the consequences of the slide. Analysts can alter the marking of a Petri Net to indicate the different conditions that hold for operators and their systems. These walk-throughs can be used to simulate the sequences of events and states that arise during accident scenarios. A transition can fire if all of its input places contain at least one token. After firing, a token is deposited in each of the output places of a transition. A single token is removed from all of the input places to that transition. In Figure 9.18 it is possible for the transition indicating that the first fire inspector questions the EPAI foreman about the need to secure the gas line to fire. The transition showing that the second fire inspector also questions the EPAI foreman about the need to secure the gas line can also fire. If these transitions fired then the places indicating that the Foreman is considering there comments would be marked. The transition showing that the Foreman decides to support the gas line can only fire if both of these places were marked together with the place indicating that the gas line is still exposed.

Incidents and accidents are often caused by the interaction between many different, concurrent users and systems [80, 277]. Figure 9.18 shows how Petri Nets can be used to represent one aspect of the interaction. In particular, this diagram shows how the first and second inspectors persuade the foreman to shore the gas line with saw horses. Although Figure 9.18 does not represent the real-time characteristics of the Allentown incident, it does accurately represent more abstract synchronisation properties. For instance, both the first and the second fire inspectors must question the need to support the pipeline before the Foreman considers supporting it. This is represented in Figure 9.18 by the places that lead to the transition labelled Foreman decides to shore the gas pipe with saw horses. This transition cannot fire until both of the places are marked to show that the Foreman is considering the implications of the inspectors' warning.

Figure 9.18 illustrates the common observation that initial failures seldom lead 'directly' to safety-critical incidents. The foreman had the opportunity to avert the Allentown explosion by correctly supporting the gas line. Indeed, the actions that he took in shoring-up the pipeline may have delayed its failure. Figure 9.18 also illustrates another important point about the reconstruction of complex failures. The resulting models often embody particular views and assumptions about the events leading to an incident. For example, the NTSB investigation obtained witness statements from Housing Association employees who:

> "...frequently passed the excavation between May 23 and June 9 stated they observed that the exposed pipe was not supported." [588]

This statement is ambiguous. It is difficult to be certain whether the employees could not see the supports, whether they saw the supports and believed them to be insufficient or whether there really were no supports there at all. Figure 9.18 does not consider such additional evidence and simply shows that the saw horses were in place throughout this period. However, Petri Nets can be used to develop alternative reconstructions that reflect these different interpretations of the available evidence. If the differences between these models were considered to be significantly important to

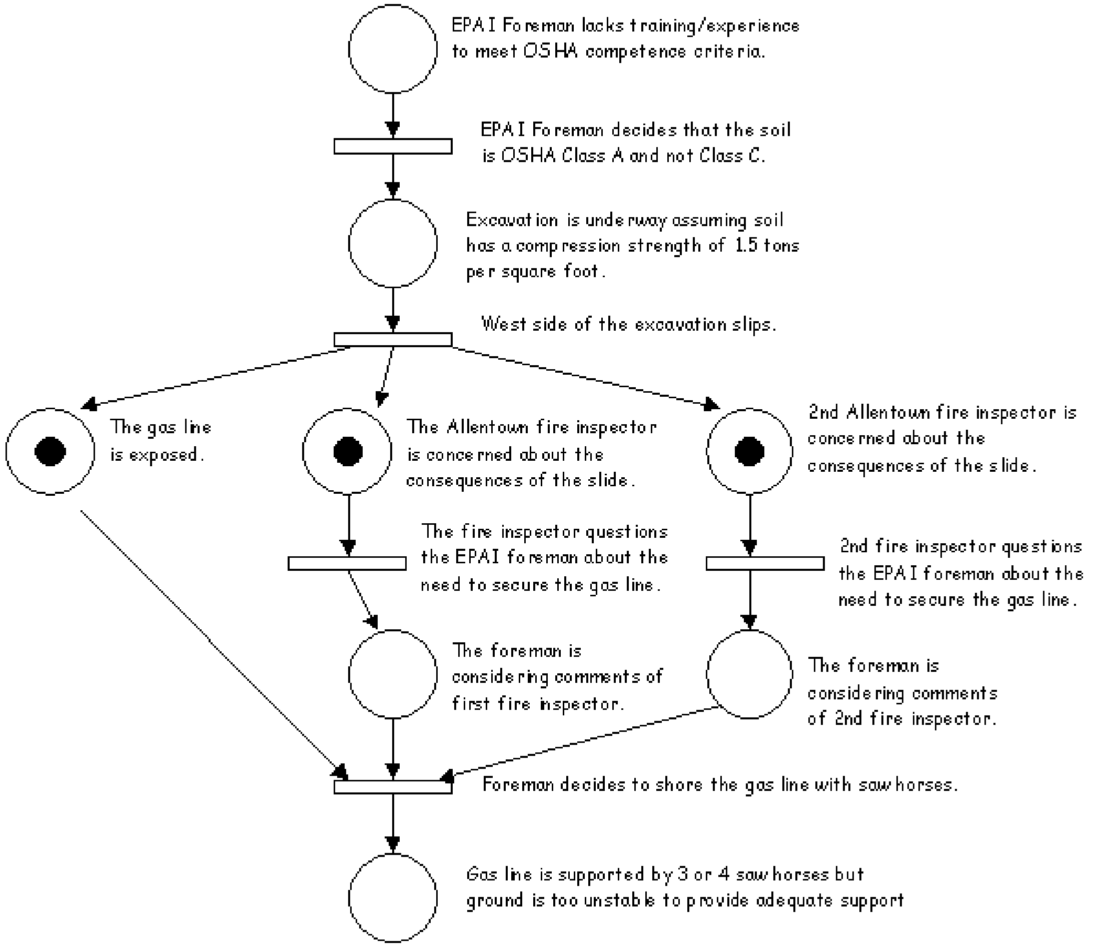
EPAI Foreman lacks training/experience to meet OSHA competence criteria.
EPAI Foreman decides that the soil is OSHA Class A and not Class C.
Excavation is underway assuming soil has a compression strength of 1.5 tons per square foot.
West side of the excavation slips.
The gas line is exposed.
The Allentown fire inspector is concerned about the consequences of the slide.
2nd Allentown fire inspector is concerned about the consequences of the slide.
The fire inspector questions the EPAI foreman about the need to secure the gas line.
2nd fire inspector questions the EPAI foreman about the need to secure the gas line.
The foreman is considering comments of first fire inspector.
The foreman is considering comments of 2nd fire inspector.
Foreman decides to shore the gas line with saw horses.
Gas line is supported by 3 or 4 saw horses but ground is too unstable to provide adequate support

Figure 9.18: A Petri Net With Multiple Tokens

any subsequent analysis then this should trigger further investigation. As we have seen, however, the supports were ultimately insufficient to protect the integrity of the pipeline. Figure 9.1.3, therefore, extends Figure 9.18 to show how the additional work, associated with removing the contaminated soil, placed undue stress on the exposed pipeline. It also shows how the Foreman's actions in attempting to shore-up the pipe with the saw horses can also, arguably, have helped to undermine a further defence. In particular, this partial remedy seems to have satisfied the concerns expressed by the inspectors. The first fire inspector's shift had ended by this point in the incident and so Figure 9.1.3 represents this important event by the transition labelled 2nd Fire Inspector and Foreman decide not to take any further action.

The upper components of the Petri Net in Figure 9.1.3 deal with the Foreman's decision to shore up the pipe in response to comments from the Allentown Fire Inspectors. The bottom right components deal with the catalytic events that stemmed from the decision to remove the concrete base and contaminated soil, which had surrounded the tank. The actions associated with the removal of this material placed the immediate stresses on the pipe that led to the failure of the compression coupling:

> "The tank was successfully removed from the excavation, and samples of soil were taken adjacent to the tank's concrete support, which remained in the excavation. The soil was to be tested to determine whether fuel had leaked from the tank and contaminated the surrounding soil. The EPAI foreman stated that before he and the other crewmembers left the site, they tried to support the pipe with saw horses, surrounded the excavation with orange plastic barrier fencing, put plastic sheeting over the excavation slopes, including the soil that lay beneath the pipe, and removed the equipment from the site... Fifteen days later, on June 9, after the EPAI received the test results, which showed that the soil around and beneath the concrete tank support had been contaminated, EPAI employees returned to remove the concrete support and contaminated soil... The backhoe (a track-mounted excavator) arrived about 12:30 p.m., and a hydraulic hammer was installed on the backhoe bucket to break up and remove the tank's concrete support. The foreman stated that he and his crewmembers removed the saw horses from beneath the pipe as the first step in removing the concrete support. He said they did not notice any movement of the pipe and did not smell any gas. The equipment operator, not the same person who had excavated the tank in May, used the backhoe to break up and remove the concrete and to excavate the fuel-contaminated soil. It took about 6 hours for the hydraulic hammer to break the concrete up. According to the EPAI employees, the impact of the hammer caused the ground to vibrate significantly. The backhoe bucket was used to remove the broken concrete and to load the pieces into a dump truck. The path of the backhoe bucket crossed over the pipe. The backhoe operator said that about 6:40 p.m. he moved the backhoe from a spot south of the excavation to one on the west. In moving it, he crossed a buried section of pipeline that was between the excavation and the north wall of Gross Towers. The odour of gas was first detected about 6:45 p.m." [588]

In Figure 9.1.3, this trigger event is represented by the transition labelled EPAI test results show the need to remove the concrete base and surrounding soil. This transition can fire because the place labelled Soil around the tank is contaminated with fuel is marked. If, however, the soil were not contaminated then this place would not have been marked and the transition could not have fired. However, as we know, the EPAI test result were positive. As a result, the associated transition can fire. This will deposit tokens in the output places that are connected to this transition in Figure 9.1.3. The new marking shows that the saw horses supports are removed to allow the access that is necessary for the work to commence. The marking will also show that a hydraulic hammer is used to break up the concrete base and that the backhoe's path crosses a buried portion of the pipeline.

It is important to note, however, that Figure 9.1.3 represents the events that led to the failure of the compression coupling. As with previous reconstructions in this chapter, it does not explicitly identify ways in which the incident could have been avoided. This illustrates an important point

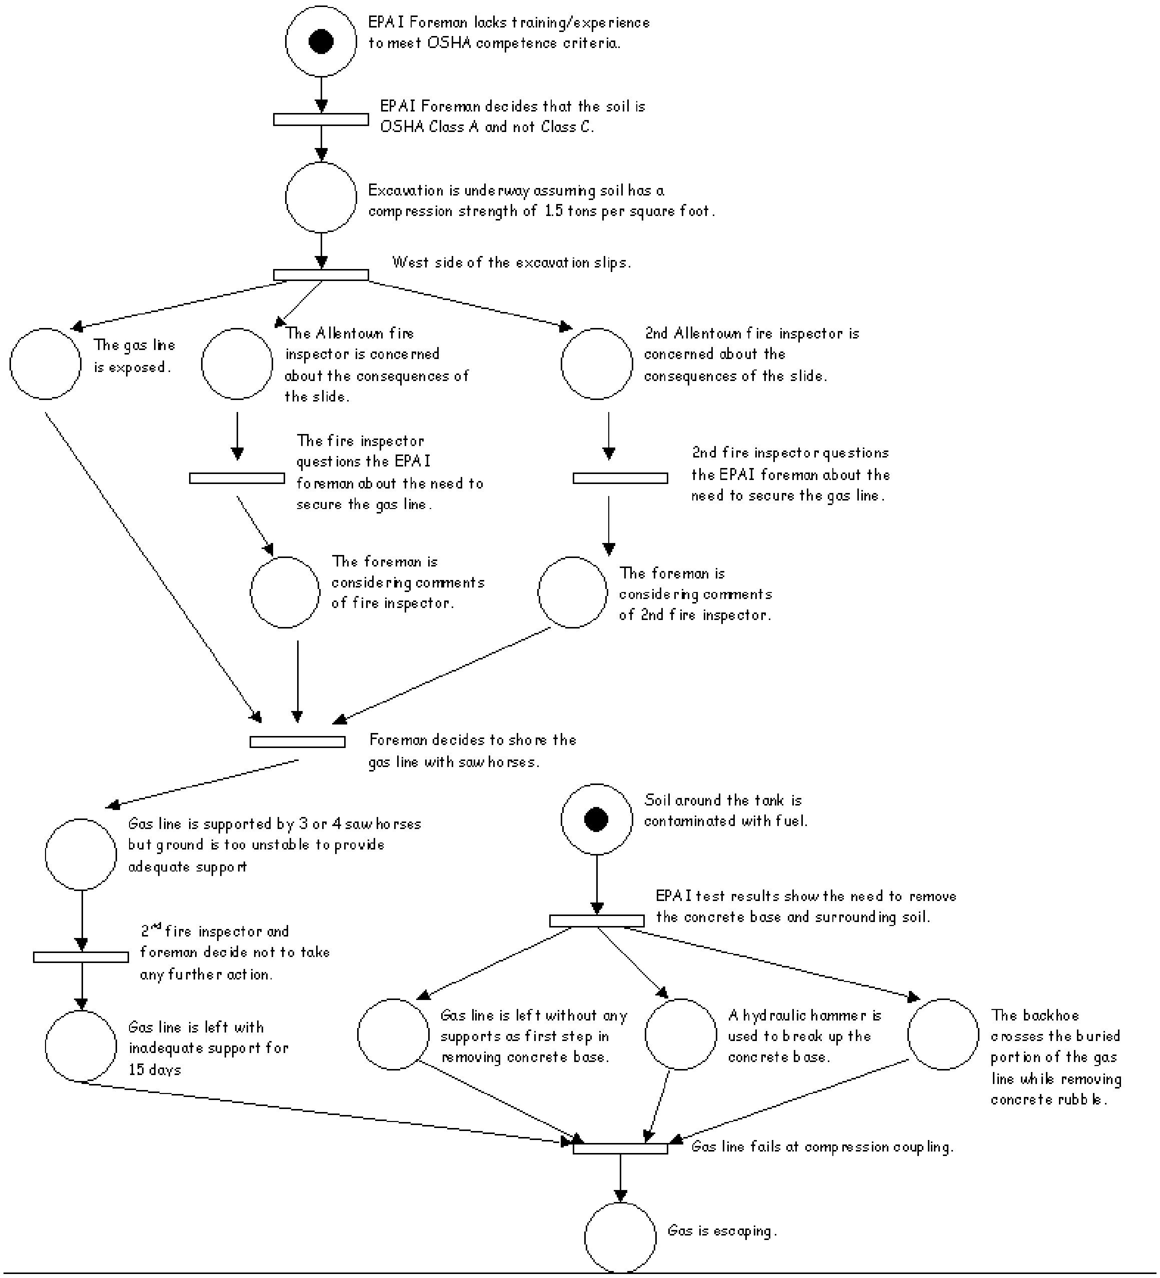
EPAI Foreman lacks training/experience to meet OSHA competence criteria.
EPAI Foreman decides that the soil is OSHA Class A and not Class C.
Excavation is underway assuming soil has a compression strength of 1.5 tons per square foot.
West side of the excavation slips.
The gas line is exposed.
The Allentown fire inspector is concerned about the consequences of the slide.
2nd Allentown fire inspector is concerned about the consequences of the slide.
The fire inspector questions the EPAI foreman about the need to secure the gas line.
2nd fire inspector questions the EPAI foreman about the need to secure the gas line.
The foreman is considering comments of fire inspector.
The foreman is considering comments of 2nd fire inspector.
Foreman decides to shore the gas line with saw horses.
Gas line is supported by 3 or 4 saw horses but ground is too unstable to provide adequate support
Soil around the tank is contaminated with fuel.
2^{nd} fire inspector and foreman decide not to take any further action.
EPAI test results show the need to remove the concrete base and surrounding soil.
Gas line is left with inadequate support for 15 days
Gas line is left without any supports as first step in removing concrete base.
A hydraulic hammer is used to break up the concrete base.
The backhoe crosses the buried portion of the gas line while removing concrete rubble.
Gas line fails at compression coupling.
Gas is escaping.

Figure 9.19: A Petri Net Showing Catalytic Transition.

about the use of graphical notations, including time-lines, Fault trees and Petri Nets. They provide concise means of capturing the events that lead to incidents and accidents. They provide communications tools and can be shown to the other participants in an enquiry. They do not provide a panacea for the problems of incident analysis. In particular, they do not replace the judgemental skills that must be developed by human factors and systems engineers. In our scenario, there is no automatic means of moving between the Petri Net representation and the remedies that can prevent an incident from recurring.

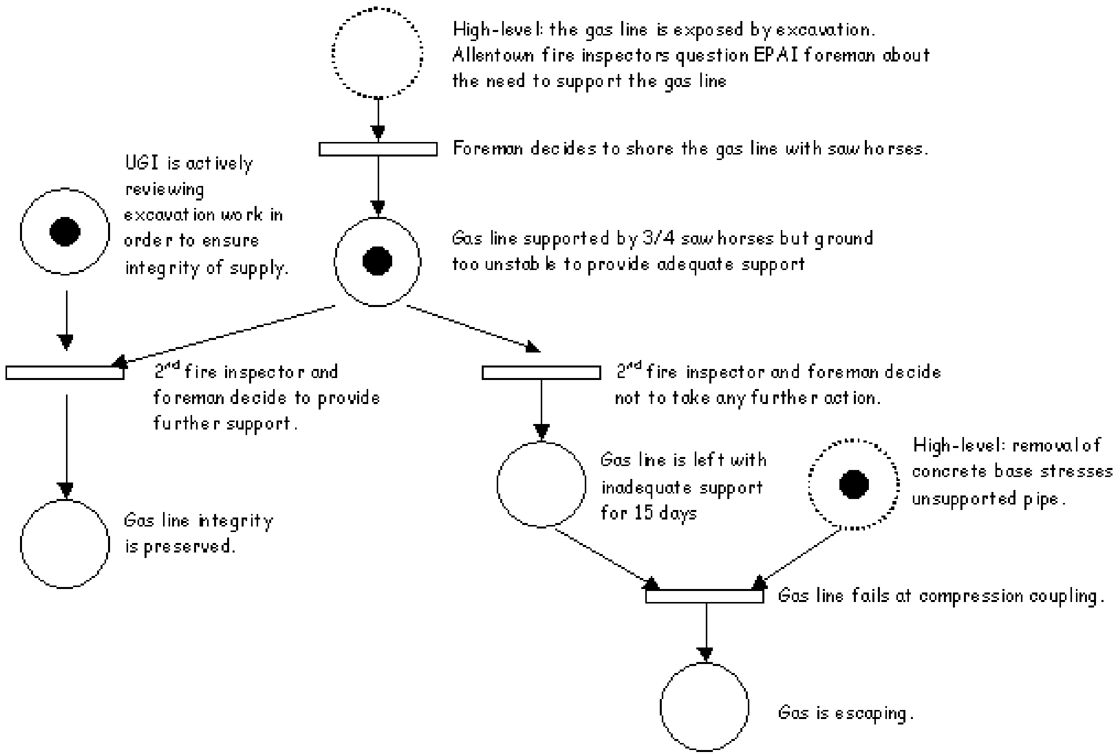
High-level: the gas line is exposed by excavation. Allentown fire inspectors question EPAI foreman about the need to support the gas line
Foreman decides to shore the gas line with saw horses.
UGI is actively reviewing excavation work in order to ensure integrity of supply.
Gas line supported by 3/4 saw horses but ground too unstable to provide adequate support
2nd fire inspector and foreman decide to provide further support.
2nd fire inspector and foreman decide not to take any further action.
Gas line is left with inadequate support for 15 days
High-level: removal of concrete base stresses unsupported pipe.
Gas line integrity is preserved.
Gas line fails at compression coupling.
Gas is escaping.

Figure 9.20: A Petri Net Showing Conflict

Previous paragraphs have shown how Petri Nets can be used to represent important events in the course of an incident. Investigators can also exploit this notation to hypothesise about alternative scenarios. Figure 9.1.3 represents two possible outcomes for the Allentown incident. One terminating place shows that gas is escaping. The other shows that the integrity of the supply is preserved. Analysts can use such networks to focus attention upon techniques that are intended to prevent future incidents. Human factors and systems engineering must be exploited so that the transition, labelled 2nd Fire Inspector and Foreman decide not to take further action, never fires. The reason we are concerned to disable this transition is that it is one possible outcome from what is known as a conflict situation. The place labelled Gas line is supported by 3 or 4 saw horses but ground is too unstable to provide adequate support is marked. As a result, it is possible to fire either the transition indicating no further action or the transition representing the decision to provide additional support. The network does not indicate which of these two possible transitions will fire. Given this marking we can, however, be sure that only one will fire and that they cannot occur simultaneously. Firing the transition indicating no further action would remove a token from the place labelled Gas line is supported by 3 or 4 horses but ground is too unstable to provide adequate support. This would disable the transition indicating that the 2nd Fire Inspector and the Foreman decide to provide further support. Conversely, firing the transition which indicates further actions would lead to a marking for the place labelled Gas line integrity is preserved. Petri Nets that include these conflict situations are non-deterministic. Any one of the transitions from a marked place can be selected for firing. In more conventional applications of the Petri Net notation it is, typically, important to detect and remove such non-determinism; it indicates an apparently random behaviour on the part of any proposed system. In incident reconstruction, however, this technique can be used to represent the non-determinism which is inherent in many complex multi-user, multi-system applications. This can, however, be problematic if investigators want to model the likely path of an incident rather than possible alternative behaviours.

Conflict situations represent critical stages in an incident reconstruction. Non-determinism indicates a loss of control over the behaviour of the 'system'. It is, therefore, important that the recommendations from an incident report will remove conflict from the Petri Net reconstruction of an incident. For example, the NTSB enquiry recommended that the excavation contractor should:

> "Modify its excavation-damage prevention program to include the review and close monitoring of any proposed excavation near a gas service line, including any line with unanchored compression couplings, that is installed near a building and that, if damaged, might endanger public safety significantly. (Class II, Priority Action) (P-96- 5)" [588]

Inhibitor arks provide a means of representing the intended effect of such recommendations. Transitions which are linked by an inhibitor can only fire if the place from which the inhibitor comes is not marked. Inhibitors are represented graphically as an edge with a small empty circle on one end. In Figure 9.1.3 an inhibitor arc is shown running from the place labelled Foreman and employees trained in OSHA and company health and safety program for excavation and training to the transition marked 2nd Fire Inspector and Foreman decide not to take any further action. The input place to this inhibitor is marked. In consequence, Figure 9.1.3 can be interpreted as stating that any decision to reject further actions cannot be taken because the Foreman's training 'inhibits' him from leaving the excavation partially supported.

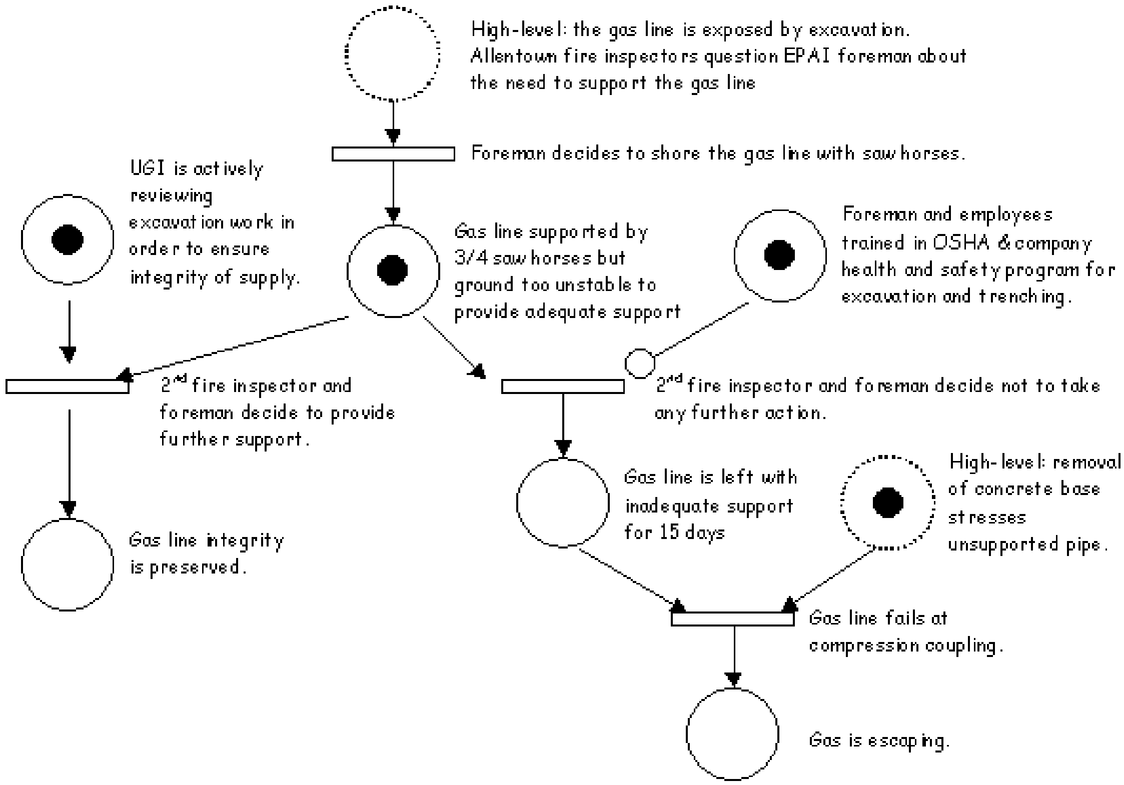

Figure 9.21: A Petri Net With An Inhibitor Avoiding Conflict.

The recommendation represented by the inhibitor arc in Figure 9.1.3 is insufficient to guarantee the safety of the system. The transition labelled 2nd Fire Inspector and Foreman decide to provide further support cannot fire unless the place marked UGI is actively reviewing excavation work in order to ensure integrity of supply is also marked. In other words, improvements in the training of excavation teams might have encouraged the foreman not to leave the gas line partially supported by the saw horses. However, this need not have guaranteed that any eventual actions would have adequately addressed the risks posed by the exposed pipeline. The participation and oversight of the gas supply company might have provided increased confidence that positive actions would be taken to address any damage that had been sustained. The place labelled UGI is actively reviewing excavation work in order to ensure integrity of supply, therefore, represents the NTSB's additional recommendation that the gas supply company must:

> "Modify its excavation-damage prevention program to include the review and close monitoring of any proposed excavation near a gas service line, including any line with unanchored compression couplings, that is installed near a building and that, if damaged, might endanger public safety significantly. (Class II, Priority Action) (P-96- 5)" [588]

In Figure 9.1.3 this place is marked and so the transition labelled 2nd Fire Inspector and Foreman decide to provide further support can fire. This in turn will mark the place indicating that Gas line integrity is preserved.

Previous Petri Nets represent the Allentown incident at an extremely high level of abstraction. This is inappropriate for the later stages of incident reconstruction. For instance, it may be necessary to model the detailed gas flow into the Housing Association's building. In fact, this was done to determine that the gas flowed underground to Gross Towers. It then passed through openings in the buildings foundation into the space beneath the mechanical room, which served as a combustion air intake reservoir for boilers. The gas then passed through openings in the floor of the building's mechanical room from where it migrated to other floors through the adjacent boiler exhaust tower, through a rubbish chute and through floor openings for electrical and other building services. It may also be important to reconstruct the more detailed cognitive and perceptual factors that influence an individual's response to potential accidents. For instance, the NTSB interviews revealed that the Foreman did not share the First Fire Inspector's concerns because he believed that the pipe did not use compression joints:

> "The fire inspector said that he questioned the EPAI foreman about the need to secure the gas line. He said that the foreman told him the condition presented no problem because the gas line was an all welded system. (The foreman later stated that based on his experience he believed all gas systems were welded)." [588]

This reconstruction is revealing because it implies that the inspector was prepared to accept the foreman's judgement. He assumed that the foreman had greater technical competence than, in fact, he did. Petri Nets can also be used to model these details. Places and transitions can be replaced by sub-networks to provide finer grained representations. The transition labelled High-level: The Fire Inspector questions the EPAI foreman about the need to secure the gas line can be refined into the sub-network shown in Figure 9.1.3.

Ths more detailed reconstruction of the incident can help to generate further hypotheses and questions. For instance, the previous paragraphs have focussed on the NTSB's recommendations about the need to improve the training of excavation crews. They have also incorporated the recommendations for improved monitoring by service suppliers into Figure 9.1.3. However, experience has shown that improved training and manual surveillance cannot be relied upon to guarantee the safety of future systems. In consequence, the NTSB investigators focussed most of their attention on the potential benefits of EFV's. These were discussed in the section on graphical time-lines. However, the Petri Net reconstruction of Figure 9.1.3 also suggests questions about the nature and origin of the disrepancy between the Foreman's mental model of the pipeline construction and the actual techniques that were used to build it. In particular, subsequent analysis might focus on why compression joints are not routinely anchored to provide increased protection against longitudinal

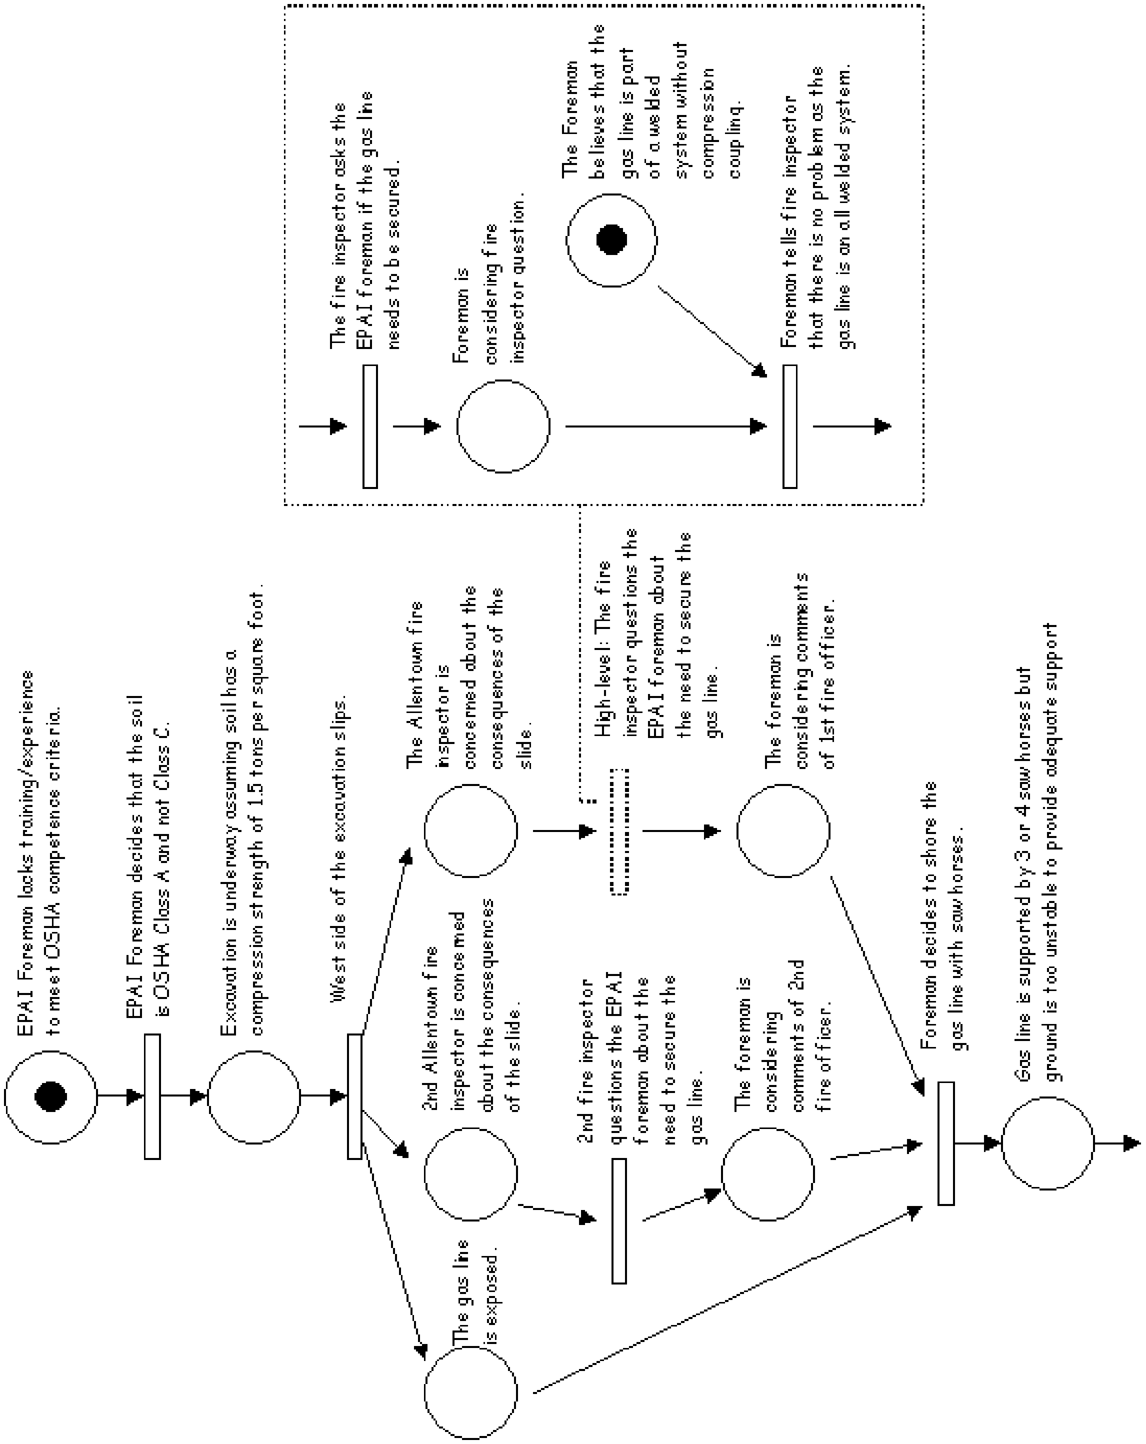
EPAI Foreman lacks training/experience to meet OSHA competence criteria.
EPAI Foreman decides that the soil is OSHA Class A and not Class C.
Excavation is underway assuming soil has a compression strength of 1.5 tons per square foot.
West side of the excavation slips.
The Allentown fire inspector is concerned about the consequences of the slide.
2nd Allentown fire inspector is concerned about the consequences of the slide.
The gas line is exposed.
High-level: The fire inspector questions the EPAI foreman about the need to secure the gas line.
2nd fire inspector questions the EPAI foreman about the need to secure the gas line.
The foreman is considering comments of 1st fire officer.
The foreman is considering comments of 2nd fire officer.
Foreman decides to shore the gas line with saw horses.
Gas line is supported by 3 or 4 saw horses but ground is too unstable to provide adequate support
The fire inspector asks the EPAI foreman if the gas line needs to be secured.
Foreman is considering fire inspector question.
The Foreman believes that the gas line is part of a welded system without compression coupling.
Foreman tells fire inspector that there is no problem as the gas line is an all welded system.

Figure 9.22: A Sub-Net Showing Crew Interaction.

pressures. The NTSB investigators considered introduced this issue but never took it any further in either their reconstruction or analysis of the incident:

> "A note on the UGI's original service record stated that the line was 'Tied in Solid,' meaning that the pipe lengths were welded. However, to comply with 1971 Federal requirements on protecting steel pipelines against corrosion, the UGI began installing corrosion-protection systems on segments of its pipeline systems that had been installed before the requirements were adopted. The UGI's records show that on September 27, 1973, an electrically insulating compression coupling 9 was installed in the service line. Although there is no documentation of the instructions given the crewmembers about the work, records and physical evidence show that they installed an insulating compression coupling in the service line north of the wall next to the boiler room. That coupling was installed just inches south of a noninsulating compression coupling for which there are no records and which was apparently installed at the same time as the insulating coupling to obtain adequate space to install the insulating coupling. Neither compression coupling was anchored or otherwise protected against movement relative to the service pipe, nor were there any requirements for doing so." [588]

Given that the Foreman believed that the pipe was of welded construction and that it had greater longitudinal strength than it actually did, it seems important to consider the reasons why he eventually decided that the line should be support. The Petri Net in Figure 9.18 shows that this was the result of the combined comments of two of Allentown's Fire Inspectors. This reconstruction emphasizes the importance of providing confirmatory advice to support a colleague's concerns about the safety of such situations. It arguably illustrates the Inspectors' success in forcing the Foreman to reconsider the situation. However, this is a flawed interpretation of the model. If the Inspectors had been sufficiently concerned then they ought to have notified the gas supplier and halter the excavation. Instead, they acquiesced in the Foreman's view that the gas line could adequately be supported by the saw horses.

Figure 9.1.3 provides an alternative view of the reason why the Foreman reconsidered his decision not to support the pipeline. His eventual decision was partly due to the intervention of the inspectors but also to a chance incident involving asphalt from the excavation:

> "The fire inspector, the EPAI crewmembers and an EPAI management representative saw a piece of asphalt paving fall about 4 feet and strike the gas pipe. The piece was large (3 by 5 feet and 3 to 4 inches thick), and the pipe was not supported. The fire inspector said that the paving permanently deflected the pipe by about a foot. He stated that before the paving hit it, the pipe was sagging, but still fairly straight." [588]

In Figure 9.1.3, the place showing that the Asphalt is close to the exposed pipeline is marked. The transition labelled Asphalt hits gas pipe can then fire. This marks a place denoting that the gas pipe is deflected by about a foot. If the place denoting the Foreman's initial judgement is also marked then the transition labelled Foreman starts to have second thought about supporting the gas pipe can fire. Clearly this reconstruction has profound safety implications; the Inspectors intervention was not sufficient to cause the Foreman to reconsider his actions. The chance event of the asphalt deflecting the pipe was, arguably just as significant. The NTSB investigators found that:

> "Because the city's fire inspectors saw on May 23 that the service line was unsupported, they could have prevented the accident. They showed proper concern about the safety of the line, especially after a piece of asphalt pavement fell on it and deformed it. However, not having been instructed to do otherwise, both inspectors relied on the EPAI foremen's assessment that the line was safe. It would have been more prudent of them to ask the pipeline owner for the assessment. The Safety Board concludes that the likely reason the fire inspectors did not tell the operator that its service line was damaged was because the inspectors did not understand the importance of notifying operators so the effects on a facility could be assessed by the operators and necessary action taken. Had the inspectors notified the UGI, it, the Safety Board believes, would have taken the necessary corrective actions, and the accident would not have happened." [588].

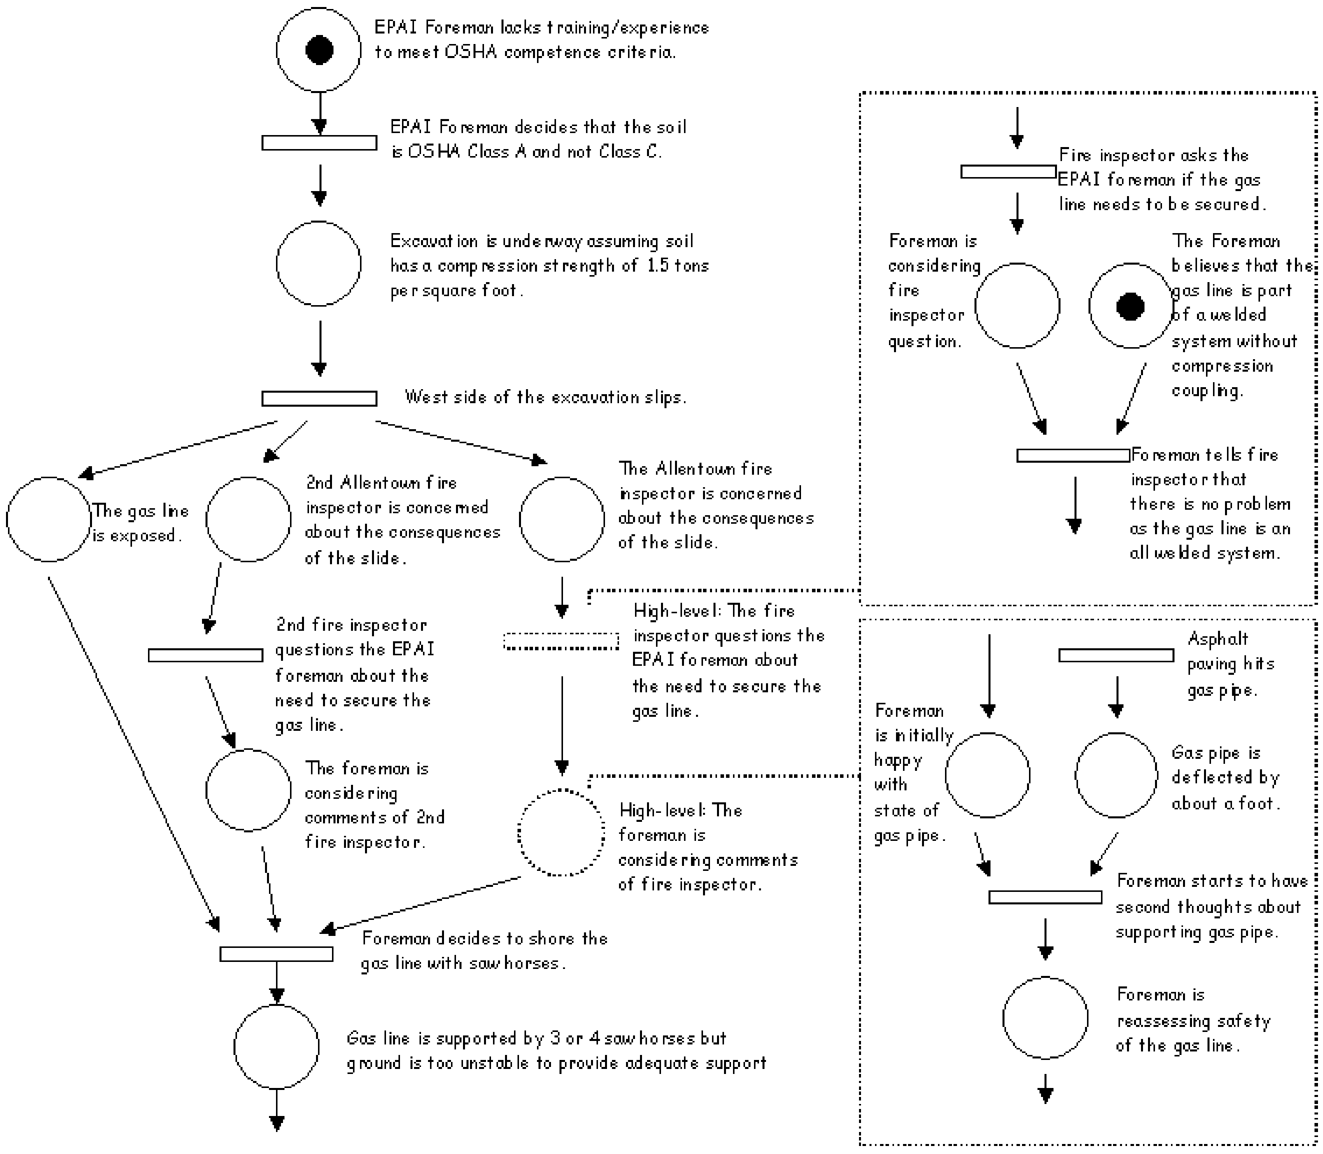
EPAI Foreman lacks training/experience to meet OSHA competence criteria.
EPAI Foreman decides that the soil is OSHA Class A and not Class C.
Excavation is underway assuming soil has a compression strength of 1.5 tons per square foot.
West side of the excavation slips.
The gas line is exposed.
2nd Allentown fire inspector is concerned about the consequences of the slide.
The Allentown fire inspector is concerned about the consequences of the slide.
2nd fire inspector questions the EPAI foreman about the need to secure the gas line.
High-level: The fire inspector questions the EPAI foreman about the need to secure the gas line.
The foreman is considering comments of 2nd fire inspector.
High-level: The foreman is considering comments of fire inspector.
Foreman decides to shore the gas line with saw horses.
Gas line is supported by 3 or 4 saw horses but ground is too unstable to provide adequate support
Fire inspector asks the EPAI foreman if the gas line needs to be secured.
Foreman is considering fire inspector question.
The Foreman believes that the gas line is part of a welded system without compression coupling.
Foreman tells fire inspector that there is no problem as the gas line is an all welded system.
Asphalt paving hits gas pipe.
Foreman is initially happy with state of gas pipe.
Gas pipe is deflected by about a foot.
Foreman starts to have second thoughts about supporting gas pipe.
Foreman is reassessing safety of the gas line.

Figure 9.23: A Sub-Net Showing Alternative Reasons for the Foreman's Decision.

Previous sections have argued that the reconstruction of an adverse occurrence forms part of an iterative process. Secondary investigations provide evidence that is used to reconstruct an incident. These reconstructions help to generate causal hypotheses. The hypotheses that emerge during the analysis of a reconstruction can force investigators to continue their search for evidence. For example, the process of using Petri Nets, such as Figure 9.1.3, to reconstruct the Allentown incident leads to further hypotheses about the reasons why the Foreman did not inform the gas supplier or provide additional support for the pipeline. In particular, the Foreman did not receive any feedback to indicate that his actions had had an adverse impact upon the pipeline. There was no smell of gas and the pipe appeared to be stable:

> "The pipe deformation caused by the asphalt pavement striking the line probably caused the pipe to be pulled out partially from the coupling because of the reduction in the effective length of the pipe. However, because there was no evidence that gas was escaping from the pipe/coupling connection before June 9, it is apparent that the activities of May 23 did not cause the pipe and coupling to separate completely." [588]

The Petri net in Figure 9.1.3 can be refined to explicitly model these observations. It is important, however, to emphasise that the successive accretion of more and more details can ultimately sacrifice the tractability of this graphical notation. Investigators must have a clear understanding of the behaviour of the incident reconstructions that are represented by a Petri Net. This task can be impaired by the additional complexity that is introduced through the use of sub-networks. It can be difficult to trace the likely passage of tokens through the many places and transitions that might used to represent the cognitive, perceptual and environmental details that contribute to a complex failure. Fortunately, this task can be eased by tools that animate the enabling and firing of transitions as tokens pass from place to place in a Petri Net. For instance, Chiola's GreatSPN can be used to view tokens as they pass through a network [164]. Investigators can record which places are marked and which transitions are enabled. The ability to play these token games greatly simplifies the development of correct models. By correct here, we mean that the model reflects the investigator's view of the incident rather than that the model correctly reflects the events leading to an incident. In contrast, this latter form of correctness depends on individual investigatory skills and on the accuracy of automated logs, mentioned in previous chapters. A further advantage of Petri Net modelling tools is that the resulting animations provide powerful means of communication. They can be shown to the many different teams that must collaborate during investigations into more serious incidents.

9.1.4 Logic

Graphical notations, such as Fault trees and Petri Nets, are not the only class of analytical tools that can be used to support incident reconstruction. A number of text based formalisms can model the events that contribute to adverse occurrences. In particular, a range of logics have been used to represent and reason about incidents and accidents [118, 412, 470]. These notations have a number of important benefits for the reconstruction of safety-critical systems:

- *formally defined syntax.* Logics, typically, have well-defined syntactic rules. These rules provide a grammar that specifies how the symbols in the logic can be combined in order to form valid sentences. These rules exist for graphical notations as well. For example, places must be connected to transitions in order to form a valid Petri Net. It would make little sense to connect a place to another place.

- *clearly defined semantics.* Logics also, typically, provide procedures for deriving the intended meaning of any sentence that obeys the syntactic rules, mentioned above. This is important because considerable confusion can arise if two different analysts can derive multiple interpretations of the same sentence. It is worth mentioned, however, that the notion of a formal semantics refers only to the information that can be directly derived or proved from the sentence itself and not to any additional, subjective judgements that might be derived from subsequent analysis.

- *proof procedures.* Logics are also supported by a set of rules that define what inferences can be made from a set of sentences. These are intended to have a close relation to the informal proof procedures that we recruit in everyday life. These 'everyday' inferences can be illustrated by the following example. If we know that 'the excavation crew at Allentown work for EPAI' and that 'the Foreman is a member of the excavation crew at Allentown' then we can conclude that the 'Foreman works for EPAI'. Proof procedures are intended to codify such inferences in order to avoid the paradoxes and fallacies that often weaken informal arguments. A paradox is a sentence that obeys the grammatical syntax rules of the language and yet is self contradictory. A good example, is the liar paradox that often frustrates the interpretation of eye witness statements. If someone says 'I am lying' then if what is said is true then it is false. If what they have said is false then it is true! Proof procedures help to identify such situations by providing rules that can demonstrate the self-contradictory nature of some grammatically valid sentences.

- *tractability.* The proof procedures, mentioned above, provide rules for manipulating the sentences of a logic to derive particular implications. There are corresponding procedures for the manipulation of textual representations for the Petri Net notation, illustrated in previous paragraphs. These techniques acknowledge that for anything but the simplest procedures it is more tractable to manipulate a textual rather than a graphical formalism. Unfortunately, it can be more difficult for non-mathematicians to interpret the meaning of textual representations than their graphical counterparts. As a result, tool support is often a necessary prerequisite for the commercial application of these techniques.

- *tool support.* Logic is an example of what have become know as 'formal methods'. These are mathematically based notations that possess the syntax, semantics and proof procedures, mentioned above. The precision and rigour provided by these features is argued to provide the increased assurance that is necessary when safety is at stake. As mentioned above, however, these benefits are often achieved at the cost of comprehension. It can be extremely difficult, even for skilled specialists, to perform the manual manipulation of mathematical sentences that are required by complex design tasks. Paradoxically, this can be an extremely error-prone activity. As a result a number of automated tools, theorem provers and model checkers, have been developed to support these tasks. Work is just beginning to improve our understanding of the errors that emerge even with this tool support [20, 21] .

- *application to both human factors and systems engineering problems.* As mentioned above, logic has been used to support the systems engineering of a range of safety-critical systems. It is also being applied to a range of human-system interaction problems. Most noticeably, a number of authors are using formal specification techniques to analyse the sources of mode confusion problems within the aviation industry [192]. Their work identifies areas in which the autopilot behaviour does not support the users' model of how the automation behaves. Ths commercial up-take of these ideas support the application of mathematically-based techniques to other forms of 'break down' during the operation of safety-critical systems.

The previous list provides some of the reasons that justify the application of logic to help reconstruct the events leading to incidents and accidents. As mentioned, however, it can initially be difficult to interpret the meaning of these formal notations. In consequence, the following pages will also provide informal readings for any notation that is presented.

Critical Components

A limitation with natural language approaches to incident reconstruction is that it can be difficult to identify critical information from a mass of background detail. For example, the NTSB's investigation into the Allentown explosion produced the following observations:

> "Post-accident surveys of 115 residents show that three Towers East occupants, in units 108, 408, and 902, had smelled gas immediately before the explosion and that two

> other occupants had smelled gas shortly before the explosion while they were in the mail room on the first floor. The occupant of unit 108 stated that he had reported the gas odour to '911,' but after the explosion." [588]

The results of this survey helped to form a more complete picture of the incident. Investigators must, however, determine whether such details are relevant to their subsequent analysis. This is important because hundreds and even thousands of items of evidence can be collected in the aftermath of a major incident. In fact, this survey were only mentioned as a parenthesis within the NTSB's final report. It might, therefore, be decided that such details could justifiably be omitted from any high-level reconstruction of the incident. The development of a logic-based model helps this process because investigators must identify significant categories of components that were involved in an adverse occurrence. The following list indicates some of the categories that have been identified from previous incidents:

- people. It is necessary to represent the people involved in an incident so that investigators can follow the way in which operator intervention affects the course of system failures;

- physical locations. It is necessary to represent the place in which an incident occurs because the location of a failure can have a profound impact upon an operator's ability to respond to that incident [404];

- warning systems. Investigators must also record the role that particular warning systems did or did not play in the course of an incident. For example, excess flow valves and gas detection equipment might have provided additional warnings about the Allentown incident;

- utterances. It is vital to represent communication between the operators that are involved in an incident. Misunderstandings have a profound impact upon the safety of many applications;

- tasks. It is necessary to identify the tasks that operators were or should have been performing during an incident if investigators are to understand the ways in which human intervention safeguarded the system or exacerbated any key failures.

For example, the following except is taken from the NTSB investigation into the Allentown incident. This quotation identifies important physical locations, such as the parking lot that the Foreman later attempted to call 911 from. It is also possible to identify key individuals, such as the Foreman, the Backhoe operator and the loader. We can also identify items of equipment such as the excavator's tools that failed to operate the valve:

> "While he was making the calls, the foreman said, he instructed the operator and the loader to trace the gas line back toward Utica Street until they found the shutoff valve. They found the valve near the north edge of the parking lot, but were unable to close it. They lacked the necessary tools to operate the below-ground valve. (Later, when the fire department representatives arrived, the EPAI workmen did not tell them they had been unable to close the valve.)" [588]

Table 9.1.4 summarises the entities that will be used in the logic-based reconstruction of the Allentown incident. It is incomplete in that the elements in the list can be expanded to enlarge the scope of the reconstruction. The identification of key individuals, locations, tasks etc is a manual, skill-based activity. It involves the subjective judgement of individual investigators. However, the outcome of this process is subject to debate and review because it can be explicitly represented in this tabular format. In formal terms, the elements of this table define the types that model the Allentown incident. The process of building such a table helps to strip out 'irrelevant' detail that can obscure critical properties of any reconstruction.

Axiom for the Accident System

The identification of people, physical locations, communication systems, equipment, utterances and tasks is of little benefit if analysts cannot represent and reason about the manner in which these components influence the course of an incident. The following section uses a simple form of temporal logic to demonstrate how this might be done for the Allentown case study.

People/Agents	Physical Locations	Warning Systems
backhoe_operator	utica_parking_lot	gas_detector
foreman	gross_towers_valve	
ugi	gross_towers	
fire_dept		
loader		
answer_service		
house_engineer		
housing_authority		
residents		
third_floor_resident		

Utterances/Messages	Tasks
gas_leak	initiate_evacuation
gas_line_hit	ventilate_building
trace_to_valve	shut_off_gas

Table 9.1: Critical Entity Table for the Allentown Incident

Operators and Locations

It is important to consider the physical location of system operators during major incidents. For instance, it is important to trace the movements of the foreman and the excavation crew after the gas was detected and before the explosion because it their locations provide valuable insights into their response to the incident. As we shall see, an appropriate response would have been to send workers to evacuate Gross Towers. Instead, the foremen sent his workers to turn off the gas supply with tools that could not achieve this goal. We can reconstruct the movements of these individuals from witness testimonies and the observations of NTSB investigators:

> "The foreman said that he then went to his pickup truck and, using his cellular phone,2 called the gas company and the housing authority, telling them that he was excavating near the gas line and smelled gas. He stated that he next made three attempts to phone 911. He said that each time he called, there was no answer. He said he then moved his truck to another spot in the parking lot in case the phone signal to his cellular phone was being blocked. He said that at the new location he again tried unsuccessfully to call 911."

It was during these telephone calls that the foreman asked the backhoe operator and the loader to trace the gas line back to Utica Street. We do not know the exact time at which the foreman made this request but the NTSB investigators suggest that it was after the first telephone call that was logged by UGI at 18:48. The backhoe operator must, therefore, have been within earshot of the foreman in order to respond to his instruction to trace the line. The operator then left the parking lot at the request of the foreman. It may be assumed that he reached the cut-off valve at some time after the request was issued, although there is no independent verification for the exact timing. The following clauses reconstruct these observations. They exploit a simple form of temporal logic in which the binary *at* operator takes a proposition and a term denoting a time such that $at(p, t)$ is true if and only if p is true at t. The existential, ∃ quantifier (read as 'there exists') can be used to capture the uncertainty about the timing of the operators movements. The first clause states that the backhoe operator is at the Utica Street parking lot at 18:48. The second clause states that at some time, t, after 18:48, the backhoe operator is not at the gas valve for Gross Towers:

$$at(position(backhoe\_operator, utica\_parking\_lot), 1848). \tag{9.2}$$

$$\begin{aligned}\exists\, t : & \; at(position(backhoe\_operator, gross\_towers\_valve), t) \wedge \\ & \; after(1848, t). \end{aligned} \qquad (9.3)$$

A number of technical problems surround the general application of this simple extension to propositional logic. In particular, the philosophical issue of reification forces analysts to clearly state the relationship between particular terms and objects over time. This theoretical problem is less of an issue for our purposes because we are always referring to definite entities at specific times during an accident. We, therefore, retain this simple temporal framework rather than the more elaborate temporal languages in our previous work [402, 427].

It might appear that such clauses add little to the information that is provided in the prose accounts of eye witness testimonies. The process of constructing such representations does, however, encourage investigators to re-examine all of the evidence supporting such location and timing information. To illustrate the importance of this cross-checking, the final NTSB report into the Allentown explosion states that UGI logged the first phone call at 18:48, cited on page 3 [588]. The investigators' time-line in appendix C of the report, on page 81, records the initial connection to the UGI switch board at 18:46 and the telephone call itself taking place at 18:47. By 18:48, the foreman was logged as calling the home of the vice president of his company to report the incident. The fact that such inconsistencies can be propagated into a final report reflects the importance of developing accurate reconstructions.

There are further reasons for reconstructing location information. The subsequent investigation into the Allentown incident was heavily critical of the Foreman's decision to send his crew members to shut off the valve. The NTSB inspectors argued that he should have asked them to evacuate anyone inside Gross Towers. Prompt action to safeguard the people inside the building would have mitigated the consequences of any explosion that they were ill-equipped to prevent. Further insights can be derived from the process of formalising the positional information in the clauses shown above. For instance, this reconstruction says remarkably little about the precise time at which the crew member left the Foreman. This is significant because it leaves open the possibility that the request was made shortly after 18:48. In which case, the Foreman would potentially have been left without sufficient staff to respond to an evacuation request:

> "Although it was after normal business hours, the foreman first called the UGI's Lehigh Division business office (the EPAI had not obtained and provided the foreman with the UGIs 24-hour emergency telephone number). Even after contacting the UGI, he did not say, and the UGI did not question, whether the odour of gas had been detected within the building. Had the UGI known that gas was already in the building, it probably would have told him to evacuate the occupants, which he could have done with the help of his crew and the bystanders. The UGI probably also would have notified the fire department, thus giving it more time to respond." [588]

UGI never issued the instructions to evacuate the building were never issued. Hence, the precise timings in clauses (9.2) and (9.3) are not significant for the reconstruction of the events leading to this particular incident. They are, however, significant for the wider recommendations about site evacuation procedures that may be drawn from this incident. Clearly those procedures should advise against allocating personnel before contacting the relevant supply company or the emergency services.

The previous clauses do not specify the relative position of the shut-off valve outside Gross Towers or of the Foreman's truck inside the Utica Street parking lot. Such information can be introduced by formalising a three-dimensional co-ordinate scheme [404]. This was not done because clauses (9.2) and (9.3) reflect the level of detail recorded after the incident investigation. However, such details can be represented in a logic-based notation, for example to support the analysis of tyre marks in road traffic incidents. These techniques can be directly derived from formal notations that underpin many CAD-CAM systems. This example illustrates a more general benefit of using a formal language. Logic provides an explicit representation of the level of abstraction that is considered appropriate for each stage of the reconstruction process. Investigators do not need to record the relative positions of the parking lot and the shut-off valve in order to model or represent the events leading to the explosion. Such decisions are extremely important. Too much detail and important properties of a reconstruction can become obscured by a mass of contextual information. Too little detail and it will be difficult to reconstruct the specific events that contribute to an incident. Clauses, such as (9.2) and (9.3), can be left at this high level of abstract or can be refined using the detailed coordinate systems introduced in [404]. This helps to avoid the ad hoc decisions that frequently seem to be made about the amount of location information that is included in incident reconstructions [426].

Operators and Communications

Communications problems exacerbate many major incidents. They also contribute to the emergency response and to any mitigating actions that may be performed. It is, therefore, important such utterances are explicitly represented within any reconstruction. For example, the investigation into the Allentown incident identified the following communications between the Foreman and the gas pipeline operator:

> "According to the UGIs records, the foreman's call was answered at 6:48 p.m. by UGI's Central Gas Control at Reading, Pennsylvania. According to the UGIs records, the foreman said that there was a gas leak at 1337 (Allen Street) Gross Towers in Allentown and that the gas line had been hit during digging. (The foreman acknowledged telling the UGI that he was digging near the gas line and had detected the odour of gas, but said that he did not tell the UGI that he had 'hit' the gas line.) At 6:52, the UGI received a second call, which was apparently from the foreman. The call was recorded as 'Cust [customer] just called back, said they definitely hit gas line and broke it.' The UGI's procedures did not require Gas Control to notify the Allentown fire department or any other emergency-response agency of either report about the release of gas because the caller did not indicate there was an imminent threat; consequently the fire department was not called." [588]

The following clauses reconstruct aspects of this quotation.

$$at(message(foreman, ugi, gas\_leak), 1848). \tag{9.4}$$

$$at(message(foreman, ugi, gas\_line\_hit), 1852). \tag{9.5}$$

An important benefit of temporal logic notations is that analysts can go beyond the previous clauses to specify persistent properties of incident reconstructions. For example, the $\forall$ (read as 'for all') quantifier can be used to specify that at no time did UGI pass on the foreman's messages to the Fire Department. $\neg$ stands for negation. The first of the following clauses can, therefore, be read as stating that at all times during the incident, UGI did not tell the Fire Department that there was a gas leak at Gross Towers. The seconds clauses states that at all times during the incident, UGI did not tell the Fire Department that the foreman had hit a gas line:

$$\forall\, t : \neg\ at(message(ugi, fire\_dept, gas\_leak), t). \tag{9.6}$$

$$\forall\, t : \neg\ at(message(ugi, fire\_dept, gas\_line\_hit), t). \tag{9.7}$$

Similar techniques can be used to reconstruct events for which the precise time is not known. For example, we do not know the exact time when the Foreman told the operator of the Backhoe and the loader to trace the gas line back to the shut off valve. The first of the following clauses states that there exists some time, t, when the foreman told the backhoe operator to trace back the gas line to the shut-off valve. The second clauses states that there exists some time, t, when the foreman told the loader to trace back the gas line to the shut-off valve:

$$\exists\, t : at(message(foreman, backhoe\_operator, trace\_to\_valve), t). \tag{9.8}$$

$$\exists\, t : at(message(foreman, loader, trace\_to\_valve), t). \tag{9.9}$$

Additional clauses can be introduced to narrow down the time when such an order could have been given. For instance, the investigators' statements record that it was issued *while* the foreman was making the phonecalls. The initial call to UGI was made at 18:48. Additional evidence must be found to identify the timing of foreman's final call by which time the order must have been given:

> "According to the housing authoritys records, the foreman called the housing authority at 6:55 and was connected to the after-hours answering service. The answering services records show that the foreman advised that 'they [the EPAI] were digging and they think they got the gas line.' At 7:06, according to the answering service, the foreman's message was relayed to one of the housing authority's maintenance employees, who promptly went to Gross Towers. The records of both the UGI and the housing authority of the foremans calls do not show that he said anything about detecting a strong odour of gas within the building." [588]

The following clause, therefore, states that the order to trace the gas line to the shut-off valve was made between the start of the first UGI call at 18:48 and the end of the call to the Housing Association at 19:06:

$$\begin{aligned}\exists\, t : &\ at(message(foreman, backhoe\_operator, trace\_to\_valve), t) \wedge \\ &\ after(1848, t) \wedge after(t, 1906). \end{aligned} \tag{9.10}$$

It is possible to impose stricter timing constraints than those shown in the previous clause because we know that the first explosion occurred at 18:58. It seems likely that the foreman directed his men to isolate the supply before the explosion. However, this is not explicitly indicated in the NTSB reconstruction which simply notes that the request was made at "6:??pm". These same logic-based techniques can be used to reconstruct more complex verbal exchanges, such as the transfer of messages between the Foreman, the Housing Authority answering service and the maintenance employee:

$$\begin{aligned}\exists\, t : &\ at(message(foreman, answer\_service, gas\_leak, 1855) \wedge \\ &\ at(message(answer\_service, house\_engineer, gas\_leak), t) \wedge \\ &\ after(1855, t) \end{aligned} \tag{9.11}$$

It is important to note that the preceding clauses do not represent the precise verbal components of each utterance. This information could be introduced if it were available, for instance through studying cockpit voice recordings in the aviation and shipping domains. In the case of the Allentown incident there was no such record. We only have the second-hand account of the answering service that the foreman had said "they [the EPAI] were digging and they think they got the gas line". After the incident, the Foreman denied saying that the backhoe had actually hit the line. However, the housing authority and UGI employees believed that this had been stated in his calls to them. Place holders, such as *gas_leak*, are used to capture the recollected sense of the communication without specifying its exact form.

Reasoning About Incidents

The previous section focussed on the flow of communication between the individuals and groups who were involved in an incident. This enables analysts to trace the way in which operators helped to exacerbate or mitigate the consequences of an incident. The same techniques can also be used to represent and reason more narrowly about the failure of particular system components. For instance, the emergency lighting failed during the Allentown incident:

> "Gross Towers, like all other housing complexes operated by the housing authority , had an internal fire alarm system that had alarm bells on each floor. When the system was activated, the company that monitored it promptly called the Allentown Communications Center. Gross Towers had a gas-powered emergency generator that started automatically whenever the flow of electricity to the building was interrupted. As long as the buildings gas supply was uninterrupted, the generator provided emergency lighting in the stair wells and exit lights. During this emergency, however, the generator did not operate because th e gas supply had been interrupted when the service line separated." [588]

This illustrates the point made in Chapter 3 that many incidents involve complex dependent system failures. The explosion that damaged the electrical power supply was caused by a gas leak that, in turn, prevented the emergency generators from working:

$$\neg\ at(electricity\_supply(gross\_towers), 1858). \tag{9.12}$$

$$\neg\ at(gas\_supply(gross\_towers), 1858). \tag{9.13}$$

$$\begin{aligned}\forall\, t : &\neg\ at(electricity\_supply(gross\_towers), t) \wedge \\ &at(gas\_supply(gross\_towers), t) \Rightarrow \\ &at(emergency\_lighting(gross\_towers, t).\end{aligned} \tag{9.14}$$

$$\begin{aligned}\forall\, t : &\neg\ at(gas\_supply(gross\_towers), t) \wedge \\ &\neg\ at(electricity\_supply(gross\_towers), t)) \Rightarrow \\ &\neg\ at(emergency\_lighting(gross\_towers), t)\end{aligned} \tag{9.15}$$

Previous paragraphs have used temporal logic to formalise the events leading to an accident. This formalisation process helps to strip out the contextual detail that hides critical observations in the many hundreds of pages that form conventional reports. We have not, however, shown that this approach can be used to reason about the events that lead to an incident. Rules of inference can be used to direct reasoning about an incident reconstruction. These rules are intended to increase the precision and rigour that is used when investigators draw particular conclusions from the events that they model. The general idea behind logical proof can be illustrated by the simple example that was presented in the previous paragraph. This provided a number of implications. For example, it was stated that the emergency lighting comes on if the electricity supply has failed but the gas supply is still working. It was also stated that if the gas system has failed then the emergency lighting would fail as well. We can use these assertions to make several inference if we have a proof rule of the following form. This states that if we know that some formula p is true at all times t and we know that if p is true at t then q is true at t then given we already know p is true then we can safety conclude that q is true at t as well:

$$\forall\, t, p(t), p(t) \Rightarrow q(t) \vdash q(t) \tag{9.16}$$

Given this rule we can begin to construct a formal proof to show that the emergency lighting failed in our reconstruction as a logical consequence of the gas leak. The proof begins by instantiating the particular moment of failure into the clauses introduced in the previous section:

$$\begin{aligned}&\neg\ at(gas\_supply(gross\_towers), 1858) \wedge \\ &\quad \neg\ at(electricity\_supply(gross\_towers), 1858) \Rightarrow \\ &\quad \neg\ at(emergency\_lighting(gross\_towers, 1858) \\ &\qquad \textit{Instantiate } t \textit{ in } (9.15) \textit{ with } 1858\end{aligned} \tag{9.17}$$

Given the previous proof rule and the fact that we know from clause (9.13) that the gas and electricity did fail at 18:58, it can now be concluded that the emergency lighting did not come on at that time.

$$\begin{aligned}&\neg\ at(emergency\_lighting(gross\_towers, 1859) \\ &\qquad \textit{Application of } (9.16) \textit{ to } (9.17) \textit{ given } (9.13) \textit{ and } (9.12).\end{aligned} \tag{9.18}$$

We might like to argue that there was some time after the explosion when there was still a sufficient supply within the emergency generators to drive the emergency lighting. The same procedures cannot, however, be used to prove this. Recall that clause (9.14) specified that the emergency lights came on if the electricity failed and the gas system was functioning. We know from (9.12) that the electricity failed at 18:58. However, we cannot prove from our reconstruction that the gas system was functioning at 18:58. Hence we cannot apply rule (9.16). If an investigator wished to establish that the generators were able to function for some initial time then additional evidence would have to be found. This might then support the following inference:

$$\begin{array}{l} at(gas\_supply(gross\_towers), 1858)). \\ Assumption \end{array} \tag{9.19}$$

$$\begin{array}{l} \neg\ at(electricity\_supply(gross\_towers), 1858) \wedge \\ at(gas\_supply(gross\_towers), 1858) \Rightarrow \\ \quad at(emergency\_lighting(gross\_towers, 1858) \\ \qquad Instantiate\ t\ in\ (9.14)\ with\ 1858 \end{array} \tag{9.20}$$

$$\begin{array}{l} at(emergency\_lighting(gross\_towers, 1858) \\ \qquad Application\ of\ (9.16)\ to\ (9.20)\ given\ (9.12)\ and\ (9.19) \end{array} \tag{9.21}$$

The Allentown investigators argued that:

> "Once the line and coupling separated, the EPAI could have limited the consequences. When the EPAI foreman was told about the strong odour of gas within the building, he should have immediately called 911. Contrary to his post-accident statement, telephone records show that he did not attempt to call 911 until after the explosion. Had he immediately reported the emergency to the fire department, it would have known almost 15 minutes before the explosion, giving it enough time to respond, notify the UGI, initiate evacuations and building ventilation, and, using the UGI responders, shut off the flow of gas into the building, which would have either prevented the explosion or reduced its force. The Safety Board concludes that the consequences of this accident could have been significantly reduced had the foreman promptly called 911 and had his helper promptly told the occupants of the building to evacuate." [588]

It is possible to use this statement together with the timing information that was provided in an NTSB inspector's time-line to reconstruct a number of important observations about the Allentown incident. The smell of gas was first reported by an EPAI employee to the foreman at 18:45. A statement from a passing policeman recorded the time of the explosion at 18:58. Between these two times, the foreman managed to call both UGI and the Housing Association but did not succeed in reaching the emergency services on 911.

$$\begin{array}{l} \exists\, t, t', \forall\, t'' : \\ \quad at(message(foreman, ugi, gas\_leak), t) \wedge \\ \quad at(message(foreman, housing\_authority, gas\_leak), t') \wedge \\ \quad \neg\ message(foreman, fire\_dept, gas\_leak), t'') \wedge \\ \quad before(1845, t) \wedge before(1845, t') \wedge before(1845, t'') \wedge \\ \quad before(t, 1858) \wedge before(t', 1858) \wedge before(t'', 1858) \end{array} \tag{9.22}$$

The term 'task' is typically used in the human-computer interaction literature to describe a collection of activities that are intended to achieve particular goals. Chapter 3 has argued that many incidents occur because individuals fail to perform particular tasks or because they select tasks whose goals are inappropriate for the context in which they are performed. It is, therefore, important that reconstructions trace the manner in which different tasks are allocated or imposed by the flow of information during an incident. Had the foreman completed a 911 call to the emergency services then the Fire Department would have been informed of the need to evacuate the building. Logic can be used to model the way in which such communications notify other people of the tasks they must perform. This could, equally, be done by using a conventional task analysis technique from the human factors literature, such as task analysis for knowledge description (TAKD) [428]. Later sections will, however, argue that formal reasoning techniques provide additional means of proving properties of incident reconstructions.

The previous quotation also stressed that had the Fire Department been notified by the Foreman then they, in turn, would have contacted UGI. Their responders would then have had time to 'shut off the flow of gas into the building, which would have either prevented the explosion or reduced its force'. This assertion can be modelled as follows. It should be noted that unlike the previous clause we do not bind the timing for t and t' to particular intervals. It is assumed that UGI shut off the supply whenever they are notified of a gas leak by the Fire Department. The *perform* predicate is used to represent an individual or group's attempt to achieve a particular task at a particular time:

$$\begin{aligned}\forall\, t : & \; \mathit{at}(\mathit{message}(\mathit{fire\_dept}, \mathit{ugi}, \mathit{gas\_leak}), t) \Rightarrow \\ & \exists\, t' : \mathit{perform}(\mathit{ugi}, \mathit{shut\_off\_gas}), t') \wedge \mathit{after}(t, t'). \end{aligned} \tag{9.23}$$

Incidents often act as a catalyst that provokes investigators to hypothesise about the introduction of particular pieces of equipment. Such alternative scenarios introduce a certain amount of additional complexity into the reconstruction process. Analysts and investigators must keep track of which clauses are being used to model any particular scenario. In particular, a contradiction would occur if clauses were introduced to simultaneously denote that gas detection equipment did and did not generate a warning. Brevity prevents a more detailed introduction to this issue, however, Burns' recent thesis identifies many of the technical problems that can arise from this aspect of formal reconstruction [118]. With these caveats in mind, it is possible to formalise alternative scenarios such as those suggested by the NTSB investigators in the previous quotation. It is important to repeat that these formalisations model or reconstruct certain aspects of an adverse occurrence. They do not capture every aspect of the prose descriptions produced by investigators, just as those prose descriptions to not capture every event that occurred during the incident itself. For example, the previous clauses do not capture the idea that had UGI and the Fire Department intervened, in the manner described above, then the explosion would either have been avoided or its energy reduced. Such notions can be formalised as properties of possible future states of the system using modal logics [118]. Such notations have the same foundations as the causal logics exploited by Ladkin's accident analysis techniques [469]. These notations provide elegant means of distinguishing between, for example, degrees of risk or notions of cause from notions of time. However, these approaches greatly increase the degree of mathematical sophistication that is necessary to reconstruct an incident. McDermid summarises many of the issues that are raised by the use of these techniques when he argues that increased expressiveness is often sacrificed at the cost of tractability and complexity [527].

The entities that were identified in Table 9.1.4 are generic in the sense that operators, tasks, utterances, physical locations etc. are central to a wide range of incidents reports [408, 426, 427]. This does not mean that the list is exhaustive. Some incidents require new types of entities to be introduced in order to model important aspects of an adverse occurrence. The significance of individual entities will also vary from incident to incident. For example, automated systems played a relatively minor role in the Allentown incident:

> "...the consequences of the accident might have been significantly reduced had the room in which the service line entered the building had a gas detector capable of alerting the occupants and the fire department. Had there been a gas detector in the room in which the service line entered, the occupants of the building and the fire department would have had 15 extra minutes in which to react. The fire department would have had time to communicate with the UGI, which might have been able to close the gas line valve soon after the separation occurred, thus preventing the accident. More likely, the accident would have happened, but much less gas would have been available to fuel the explosion, which might have substantially reduced the number of casualties and extent of the damage... contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building."[588]

Such findings create a number of problems for organisations that must prevent the recurrence of future accidents. It does not explain the impact that such devices might have had upon the course of the Allentown explosion. This ambiguity has serious consequences. Different readers might form very different conclusions about whether or not such systems would have had a significant impact upon the course of the incident [844, 699]. Formal proof techniques can be used to reason about the impact that such findings might have for any reconstruction. For instance, a gas detector warning might have prompted the evacuation of the building. The following clause does not specify that a gas leak must actually have occurred in order for an evacuation to be initiated:

$$\begin{aligned}&\forall\, t : at(message(gas\_detector, residents, gas\_leak), t) \Rightarrow \\ &\quad at(perform(residents, initiate\_evacuation), t)\end{aligned} \tag{9.24}$$

Formal reasoning techniques can be used to determine whether such assertions are supported by the evidence from a reconstruction. We can use the laws of our logic system to determine whether or not such a warning would actually have prompted the residents to leave the building. One way of doing this is to look for a situation that contradicts the previous assertion. This involves looking through the clauses of our model or reconstruction to find evidence of a situation in which the residents failed to evacuate their building in spite of a warning about the presence of gas. Ideally, the detection equipment should have identified the presence of gas almost immediately after the line had separated from the coupling at 18:45:

$$\begin{aligned}&at(message(gas\_detector, residents, gas\_leak), 1845) \Rightarrow \\ &\quad at(perform(residents, initiate\_evacuation), 1845). \\ &\quad \textit{Instantiate } 1845 \textit{ for } t \textit{ in } (9.24)\end{aligned} \tag{9.25}$$

$$\begin{aligned}&\neg\ at(message(gas\_detector, residents, gas\_leak), 1845) \vee \\ &at(perform(residents, initiate\_evacuation), 1845) \\ &\textit{Implication Law applied to } (9.25)\end{aligned} \tag{9.26}$$

Looking at the first part of this disjunction, we know that the residents did not initiate any evacuation.

$$\forall\, t : \neg\ \ at(perform(residents, initiate\_evacuation), t) \tag{9.27}$$

A passing police officer started clearing the building after he had heard the sound of the first explosion after 1858. We, therefore, have a contradiction with part of the previous clause:

$$\begin{aligned}&at(perform(residents, initiate\_evacuation), 1845) \\ &\textit{Assumption from } (9.26)\end{aligned} \tag{9.28}$$

$$\begin{aligned}&\neg\ at(perform(residents, initiate\_evacuation), 1845) \\ &\textit{Instantiate } 1845 \textit{ for } t \textit{ in } (9.27)\end{aligned} \tag{9.29}$$

$$\begin{aligned}&\neg\ at(perform(residents, initiate\_evacuation), 1845) \wedge \\ &at(perform(residents, initiate\_evacuation), 1845) \\ &\wedge \textit{ Introduction for } (9.28) \textit{ and } (9.29)\end{aligned} \tag{9.30}$$

As mentioned, formal reasoning is being used to reconstruct a situation that contradicts previous assertions about the potential role of gas detection equipment. The residents did not initiate an evacuation at 18:45. In order to derive the necessary contradiction we must also show that they were alerted to the presence of gas at this time. We know from the NTSB report that several of the residents had smelt gas by 18:45, almost immediately after the line had separated from the coupling. No evacuation was started. They only called 911 after the first explosion had occurred:

> “Post-accident surveys of 115 residents show that three Towers East occupants, in units 108, 408, and 902, had smelled gas immediately before the explosion and that two other occupants had smelled gas shortly before the explosion while they were in the mail room on the first floor. The occupant of unit 108 stated that he had reported the gas odour to ‘911,’ but after the explosion.” [588]

An EPAI employee is recorded on page 81 of the report as stating that a woman on the third floor shouted that she smelled a “heavy odour of gas’ at 18:45. It is not possible to resolve this reference against the room numbers mentioned in the previous citation. We do, however, know that this person did try to alert the other residents:

$$at(message(third\_floor\_resident, residents, gas\_leak), 1845). \tag{9.31}$$

This element of the reconstruction does not support the contradiction that was initially intended. We cannot show a situation in which the residents failed to respond to a detection *system*. However, the formal modelling does emphasise that residents were not alerted by their neighbours’ warnings and that even those who smelled gas did not take immediate action to evacuate the building:

$$\begin{aligned}&\neg\ at(message(gas\_detector, residents, gas\_leak), 1845).\\&Assumption\ from\ (9.26)\end{aligned} \tag{9.32}$$

$$\begin{aligned}&at(message(third\_floor\_resident, residents, gas\_leak), 1845) \wedge\\&\neg\ at(message(gas\_detector, residents, gas\_leak), 1845).\\&\wedge\ Introduction\ for\ (9.31)\ and\ (9.32)\end{aligned} \tag{9.33}$$

The previous clause illustrates the important point that formal modelling does not provide a panacea for the problems of incident reconstruction. The same insights can also be derived by careful inspection of the evidence that is gathered during a secondary investigation. However, such formal analysis introduces a discipline and rigour that can help investigators to reassess the assumptions that might otherwise be made about the course of an incident. For instance, as the previous clauses have shown, there is no guarantee that residents will respond to either automated or human warnings. It is for this reason that most institutions, including the Gross Towers retirement home, practice fire drills. It is pertinent to ask why these procedures are cued by the detection of fire rather than the presence of gas:

> “The executive director stated that the housing authority had procedures for evacuating the occupants and that the residents practiced the routines. For example, every 6 months the fire department conducted fire inspections and drills that also tested the evacuation procedures and emphasized how important it was for the residents to respond promptly. The drills included special precautions for the elderly and handicapped; and after a drill was held, all residents participated in a critique.” [588]

The previous paragraphs used formal reasoning to drive an analysis of the NTSB’s assertion that the lack of a gas detection system exacerbated the consequences of the incident by failing to alert the residents to the potential danger. This mirrors the observation that the fire brigade could have used the additional warning to notify the gas supply company:

$$\begin{aligned}&at(message(gas\_detector, fire\_dept, gas\_leak), 1845) \Rightarrow\\&\quad \exists\, t, t', t'' : at(perform(fire\_dept, initiate\_evacuation), t) \wedge\\&\quad at(perform(fire\_dept, ventilate\_building), t') \wedge\\&\quad at(message(fire\_dept, ugi, gas\_leak)), t'') \wedge\\&\quad after(1845, t) \wedge after(1845, t') \wedge after(1845, t'') \wedge\\&\quad after(t, 1858) \wedge after(t', 1858) \wedge after(t'', 1858).\end{aligned} \tag{9.34}$$

A warning from the gas detector results in a message being sent to UGI, at t'', between the moment when the gas is detected and when the moment when the explosion actually occurred. If we assume that a gas detection system had been installed:

$$\begin{array}{l} at(message(gas\_detector, fire\_dept, gas\_leak), 1845). \\ \textit{Assumption.} \end{array} \quad (9.35)$$

$$\begin{array}{l} \exists\, t, t', t'' : at(perform(fire\_dept, initiate\_evacuation), t) \wedge \\ \quad at(perform(fire\_dept, ventilate\_building), t') \wedge \\ \quad at(message(fire\_dept, ugi, gas\_leak)), t'') \wedge \\ \quad after(1845, t) \wedge after(1845, t') \wedge after(1845, t'') \wedge \\ \quad after(t, 1858) \wedge after(t', 1858) \wedge after(t'', 1858). \\ \quad \textit{Application of Modus Ponens to } (9.34) \textit{ given } (9.35) \end{array} \quad (9.36)$$

$$\begin{array}{l} \exists\, t'' : at(message(fire\_dept, ugi, gas\_leak)), t'') \wedge \\ \quad after(1845, t'') \wedge after(t'', 1858). \\ \quad \textit{Elimination of } \wedge \textit{ from } (9.36) \end{array} \quad (9.37)$$

As before, this formalisation suggests directions for further analysis. In particular, the previous clause would be satisfied if the fire service issued a warning at any time between 18:45 and 18:58. Clearly, information at the start of this interval might have had a greater impact upon the outcome that a warning that arrived only seconds before the explosion at 18:58. The gas supply company would have had a greater opportunity to cut off the supply before it built up within Gross Towers. The following clause assumes that the message was passed to UGI at 18:46; immediately after it was received by the fire service:

$$\begin{array}{l} \forall\, t : at(message(fire\_dept, ugi, gas\_leak), t) \Rightarrow \\ \quad \exists\, t' : perform(ugi, shut\_off\_gas), t') \wedge after(t, t'). \end{array} \quad (9.23)$$

$$\begin{array}{l} at(message(fire\_dept, ugi, gas\_leak), 1846) \wedge \\ after(1845, 1846) \wedge after(1846, 1858). \\ \textit{Instantiation of } t'' \textit{ for } 1846 \textit{ in } (9.37). \end{array} \quad (9.38)$$

$$\begin{array}{l} at(message(fire\_dept, ugi, gas\_leak), 1846). \\ \textit{Elimination of } \wedge \textit{ in } (9.38). \end{array} \quad (9.39)$$

$$\begin{array}{l} \exists\, t' : perform(ugi, shut\_off\_gas), t') \wedge after(1846, t'). \\ \textit{Application of } (9.16) \textit{ to } (9.23) \textit{ given } (9.39) \end{array} \quad (9.40)$$

One means of assessing the potential benefit of such an early warning is to compare the possible impact of a warning system with what actually happened during this incident. This follows what was done by the previous proof in which we compared the impact of an automated alarm with the warning that was issued by individual residents in Gross Towers. In this case, however, we know form page 3 of the NTSB report that UGI was informed of the gas leak in a telephone call by the EPAI foreman at 18:48. We also know from page 5 of the NTSB report that the UGI operators eventually cut off the gas supply to the building at 19:15. In other words, it took approximately twenty-seven minutes for UGI employees to reach the scene of the gas leak, to trace the damaged pipe back to the Utica Street supply and then to isolate the line to Gross Towers. we can use this information to instantiate t' in (9.40) by adding the twenty-seven minute delay to the best case estimate for the fire brigade passing the gas detector's warning to UGI:

$$\begin{array}{l} \mathit{perform}(\mathit{ugi}, \mathit{shut\_off\_gas}), 1913) \wedge \mathit{after}(1846, 1913). \\ \mathit{Instantiation\ of}\ 1913\ \mathit{for}\ t'\ \mathit{in}\ (9.40). \end{array} \tag{9.41}$$

The implications of this analysis are clear. The additional time gained by an automated gas detection system would only have bought an additional two minutes during this incident. This confirms the argument put forward by the NTSB's investigators. The warning would not have provided sufficient time in order to avoid the explosion. However, it does not necessarily confirm their analysis that the additional time might have enabled respondents to mitigate the consequences of the incident. The validity of such an assertion cannot be directly assessed from the reconstruction that has been presented in this chapter. Nor can it be directly assessed from any of the evidence in the final report into this incident.

Unfortunately, mathematical analysis provides non-formalists with an extremely poor idea of the argumentation processes that support particular conclusions. It is difficult for people without some mathematical background to understand the various proof rules that are applied during our formal analysis. The consequences of this should not be underestimated. The use of a mathematical notation does not guarantee that any analysis will be free from error. Formal proof rules are simply intended to explicitly represent the mechanisms that support particular inferences. They expose the reasoning that is implicit within an informal analysis of an incident or accident. The intention is that other investigators can use those proof rules to challenge the basis for particular arguments about an adverse occurrence. However, if those proof rules cannot easily be understood by other investigators then there is little likelihood that they will be able to challenge the inferences and arguments of their peers. Automated reasoning tools provide means of increasing confidence in such proofs even when they may not be accessible to all parties in an investigation. Some initial work has applied these theorem provers and model checkers to support incident investigation [421, 412]. More work remains to be done. The insights provided by these systems must still be communicated to many different domain experts. The following pages, therefore, present techniques that have been developed to address the communications problems that affect the formal analysis of incident reports.

Conclusion, Analysis and Evidence (CAE) Diagrams

Conclusion, Analysis and Evidence (CAE) diagrams provide a high level overview of the argument that investigators construct to support the findings of an incident investigation. They build on the products of any reconstruction to support the causal reasoning that will be the focus of the next chapter. It is appropriate to briefly introduce this technique here because we have already stressed the close links between investigation, reconstruction and causal analysis. This decision is also justified by the way in which CAE diagrams illustrate the products of formal reasoning. They can be used to overcome some of the problems of communicating these reconstruction techniques to domain experts who may not have any background in mathematical logic.

Figure 9.24 presents an initial CAE diagram for the Allentown incident. The nodes of this graph are annotated with direct quotations from the NTSB investigators. As can be seen, CAE diagrams are formed around particular conclusions about the adverse occurrence. Here C1 denotes the argument made on page 48 of the NTSB incident report that the lack of a gas detector contributed to the severity of this incident. This represents a particular instance of the counterfactual arguments, mentioned in previous sections. The incident would have been less severe if a gas detector had been installed. The consequences of the failure were exacerbated because such a device had not been installed. The conclusion that forms the root of a CAE diagram is, in turn, supported by a number of lines of analysis. In this instance, A1.1 argues that a gas detector might have enabled the fire department to communicate with UGI in order to ensure a more prompt response. The line of analysis represented by A1.2 denotes the argument that a gas detector might have provided the residents with an extra fifteen minutes in which to react.

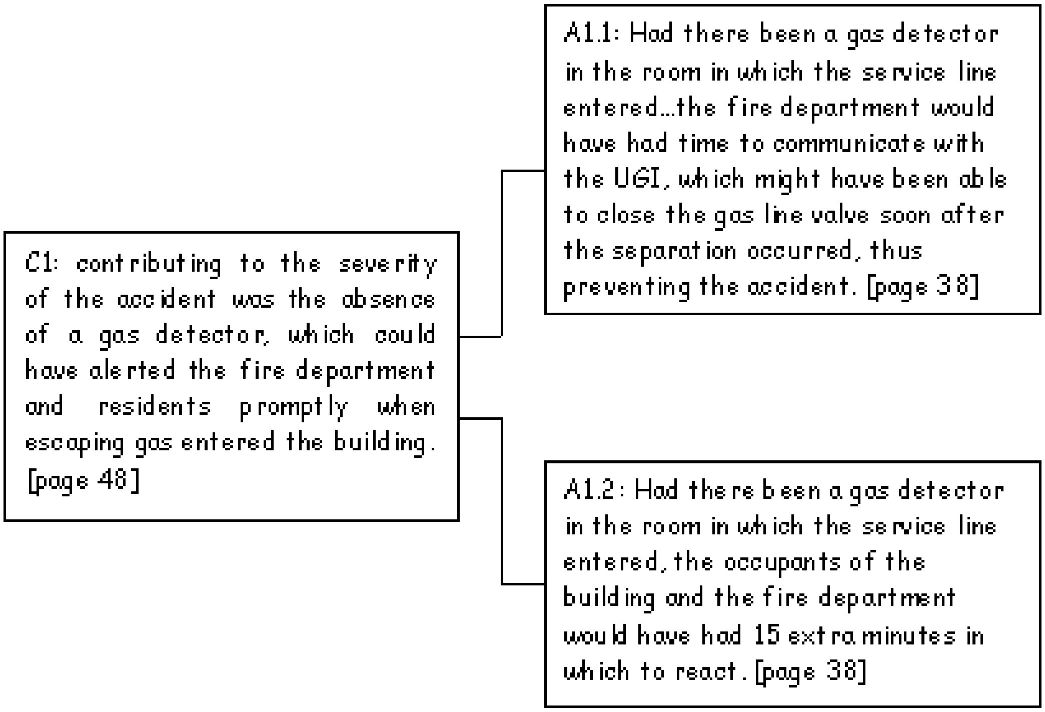
A1.1: Had there been a gas detector in the room in which the service line entered...the fire department would have had time to communicate with the UGI, which might have been able to close the gas line valve soon after the separation occurred, thus preventing the accident. [page 38]
C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]
A1.2: Had there been a gas detector in the room in which the service line entered, the occupants of the building and the fire department would have had 15 extra minutes in which to react. [page 38]

Figure 9.24: High-Level CAE Diagram for the Allentown Incident

CAE diagrams can be used to trace the arguments that both support and weaken particular conclusions. For instance, Figure 9.25 extends Figure 9.24 to show an objection to the NTSB conclusion. This is denoted by the dotted line between A1.1 and A1.1.1. Figure 9.25 counters the argument that a gas detector might have prevented the incident. The analysis in A1.1.1 argues that a gas detector would not have provided a warning soon enough for UGI to avert the explosion. This analysis is based on the assumption that it took 27 minutes to cut the supply from the time at which UGI were first notified at 18:48. Even if the gas detector had issued a warning immediately after the line was cut this could only have gained two minutes from the time at which the foreman made his first call. This line of argument is supported by two items of evidence. The node E1.1.1.1 shows that according to UGI records, the Foreman's initial call was answered at 18:48. The evidence denoted by E1.1.1.2 shows that the UGI employee only succeeded in shutting down the gas line by 19:15.

Figure 9.26 provides a further illustration of the way in which CAE diagrams sketch the arguments for an against particular conclusions. Rather than focusing on the response of the Fire Service and UGI to any automated warning, this CAE diagram illustrates a counter argument to the theory that a gas detector might have encouraged the residents to evacuate Gross Towers. This is based on the observation that some residents did know about the gas leak and yet still did not initiate an evacuation. As can be seen, two further items of evidence support this counter argument. E1.2.1.1 denotes that a resident did smell gas almost as soon as the pipeline failed. This is recorded at 18:45 on page 81 of the NTSB report. E1.2.1.2 shows that at least three other residents had first-hand knowledge of a potential gas leak but nobody rang '911' until after the first explosion. The evacuation was, in fact, initiated by a passing police officer.

Many investigators recruit extremely complex arguments both for and against particular conclusions. As can be seen, it is possible to identify a number of competing positions within the NTSB reports into the Allentown incident. CAE diagrams provide a high-level means of mapping out these positions to ensure that analysts demonstrate that their analysis is well-founded in the events that are represented within a reconstruction. This is important because there causal arguments or arguments about the mitigation of an incident can become 'detached' from the evidence that is

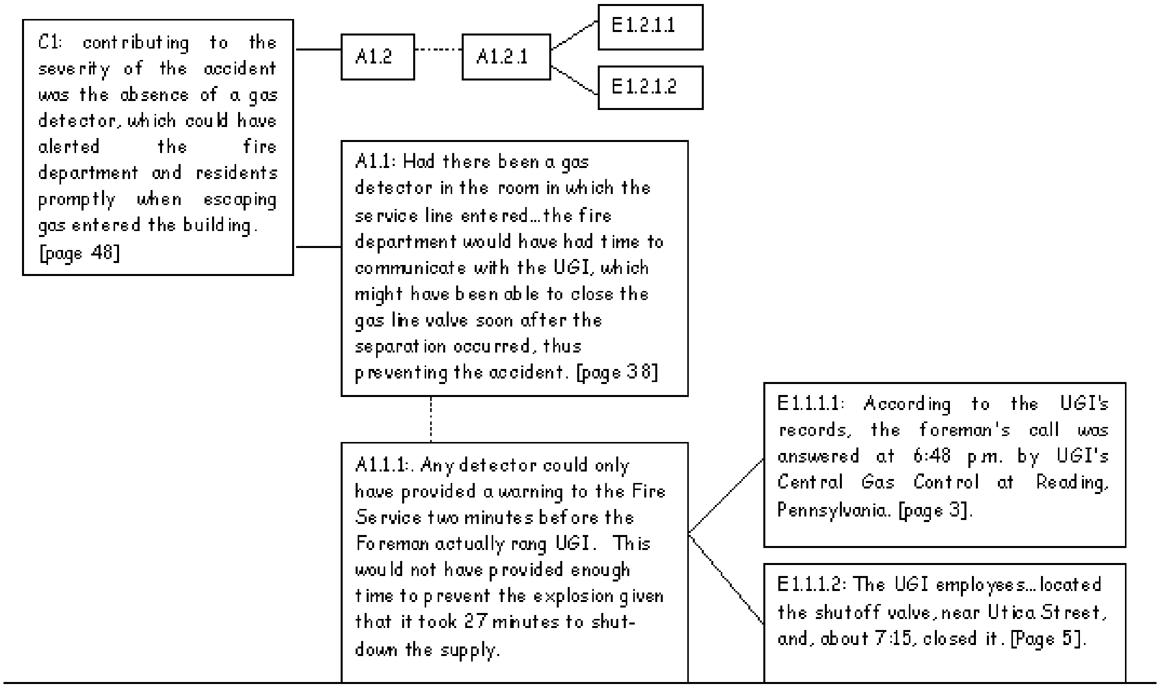
C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]
A1.2
A1.2.1
E1.2.1.1
E1.2.1.2
A1.1: Had there been a gas detector in the room in which the service line entered...the fire department would have had time to communicate with the UGI, which might have been able to close the gas line valve soon after the separation occurred, thus preventing the accident. [page 38]
A1.1.1:. Any detector could only have provided a warning to the Fire Service two minutes before the Foreman actually rang UGI. This would not have provided enough time to prevent the explosion given that it took 27 minutes to shut-down the supply.
E1.1.1.1: According to the UGI's records, the foreman's call was answered at 6:48 p.m. by UGI's Central Gas Control at Reading, Pennsylvania. [page 3].
E1.1.1.2: The UGI employees...located the shutoff valve, near Utica Street, and, about 7:15, closed it. [Page 5].

Figure 9.25: Representing Counter Arguments in a CAE Diagram (1)

gathered during a primary and secondary investigation. This need not, however, be malicious. It can simply stem from the logistical problems created by the increasing complexity of many technological failures. This is illustrated by the way in which Figure 9.26 cites evidence from page 3 to analyse arguments that were proposed on page 38 in support of a conclusion that is presented on page 48 of the NTSB report. Without such diagrammatic support, there is a danger that important evidence may be overlooked when analysing any reconstruction.

Figure 9.26 is not unusual in the complexity of the argument that it presents. For example, Figure 9.27 extends the previous analysis. It represents a line of argument that supports the assertion that a gas detector might have helped the gas supplier, UGI, to prevent the explosion. As can be seen, A.1.1.2 argues that UGI would have responded differently if a warning had been raised by the Fire Service rather than from the EPAI foreman. This line of argument is supported by two additional items of evidence. E.1.1.2.1 emphasizes the point that the foreman's calls to UGI did not emphasise the degree of threat posed by the initial gas leak. In E.1.1.2.2, UGI's records indicate that the foreman did not report the smell of gas within Gross Towers. Both items of evidence help to explain why UGI personnel might not have understood the implications of the foreman's report. The NTSB investigators argue that if the suppliers had been notified by the fire service, in response to an automated alarm, then the warning would have been less ambiguous. This would also have avoided the communications problems, noted in previous chapters, that often arise when individuals must report adverse events that they are themselves implicated in.

The previous diagrams have shown how CAE diagrams can be used to map out the arguments and counter arguments that are constructed using the evidence provided in reconstructions. This is important if analysts are to consider not simply the arguments that they favour but also the competing views that might be raised in the aftermath of an investigation. We have not, however, shown how this techniques might also be used to communicate the products of any formal reconstruction using logic or other mathematical notations. In contrast, Figure 9.28 presents a relatively simplistic means of achieving this aim. Textual annotations to the nodes in a CAE diagram are extended to include clauses derived from the formal reconstruction of an adverse occurrence. In this case, A1.1 shows that if the fire department had alerted UGI to the gas leak then they would have shut it off

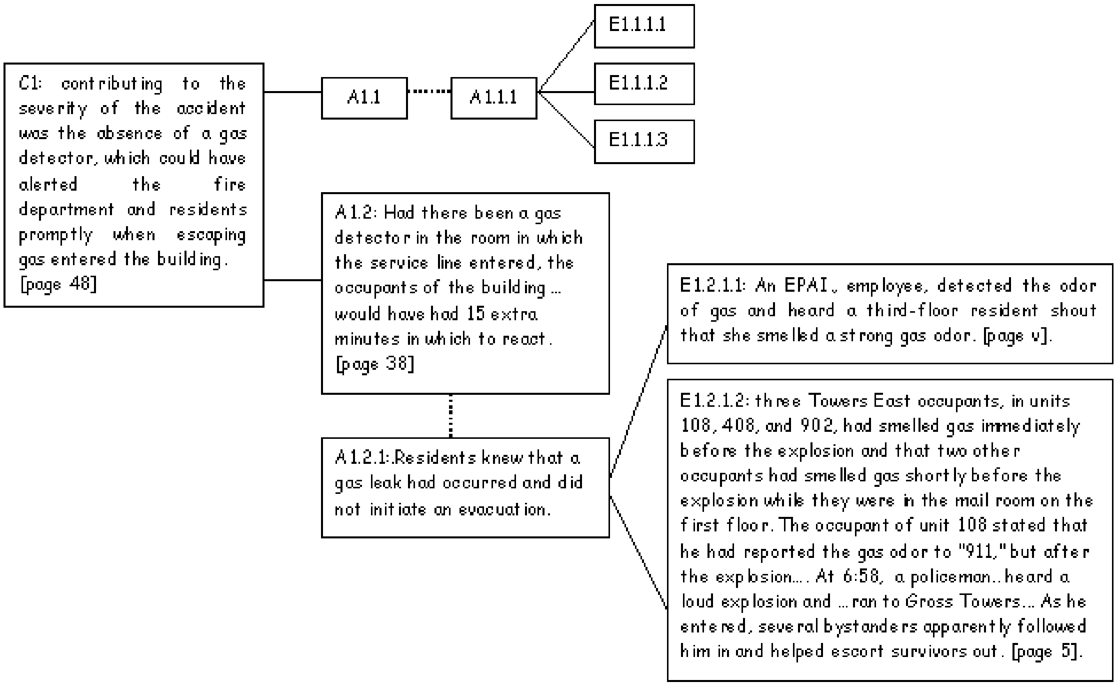
C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]
A1.1
A1.1.1
E1.1.1.1
E1.1.1.2
E1.1.1.3
A1.2: Had there been a gas detector in the room in which the service line entered, the occupants of the building ... would have had 15 extra minutes in which to react. [page 38]
A1.2.1:.Residents knew that a gas leak had occurred and did not initiate an evacuation.
E1.2.1.1: An EPAI. employee, detected the odor of gas and heard a third-floor resident shout that she smelled a strong gas odor. [page v].
E1.2.1.2: three Towers East occupants, in units 108, 408, and 902, had smelled gas immediately before the explosion and that two other occupants had smelled gas shortly before the explosion while they were in the mail room on the first floor. The occupant of unit 108 stated that he had reported the gas odor to "911," but after the explosion.... At 6:58, a policeman..heard a loud explosion and ... ran to Gross Towers... As he entered, several bystanders apparently followed him in and helped escort survivors out. [page 5].

Figure 9.26: Representing Counter Arguments in a CAE Diagram (2)

before the explosion at 18:58. A1.2 states that the residents would have initiated an evacuation if a detection system had identified the gas leak when it first started at 18:45. These formalisations represent strong requirements. For instance, A1.1 states that UGI would shut off the gas before 18:58 irrespective of the time at which the Fire Service contacted them. This seems unreallistic and, as we have seen, additional clauses may be introduced to reflect the minimum time necessary between any notification and a successful intervention by UGI. CAE diagrams, such as Figure 9.28, can help to expose such unwarranted assumptions that might otherwise be embodied within a formal analysis.

Figure 9.29 presents part of the formal reasoning that was used to assess whether or not the assumptions, embodied in Figure 9.28, might be sustained. Elements of the mathematical model constructed in the previous section are linked to the natural language evidence that was identified by the NTSB investigators. This is then used to create a conjunction which shows that the foreman alerted UGI to the gas leak at 18:48 and that their representatives did not shut the supply until 19:15, after the explosion at 18:58. As we have seen, this is not a direct contradiction of the argument put forward by the investigation team. However, it does use the evidence about what actually happened in this incident to construct a counter-case against the hypothesis about the effectiveness of an automated gas detector.

The CAE in Figure 9.30 shows how elements of the formal analysis can be used to counter the argument that a gas detector might have encouraged the resident to initiate an evacuation. Elements of the reconstruction are again linked to the natural language evidence on page v and page 5 of the investigators' report. This evidence is then used to develop a counter case. We can establish that residents did know about the gas leak almost as soon as it occurred, they smelt gas at 18:45. However, they did not initiate an evacuation in spite of this direct physical evidence of the potential danger. Human factors research into the efficacy of alarms suggests that many automated warnings have little effect on such a response, especially given that the residents had already received evacuation training [636].

It is important to emphasise that we have only shown one means of using CAE diagrams to represent the insights that can be gained from the formal reconstruction of adverse occurrences. In the previous examples, we have constructed models of the incident and then used those models to

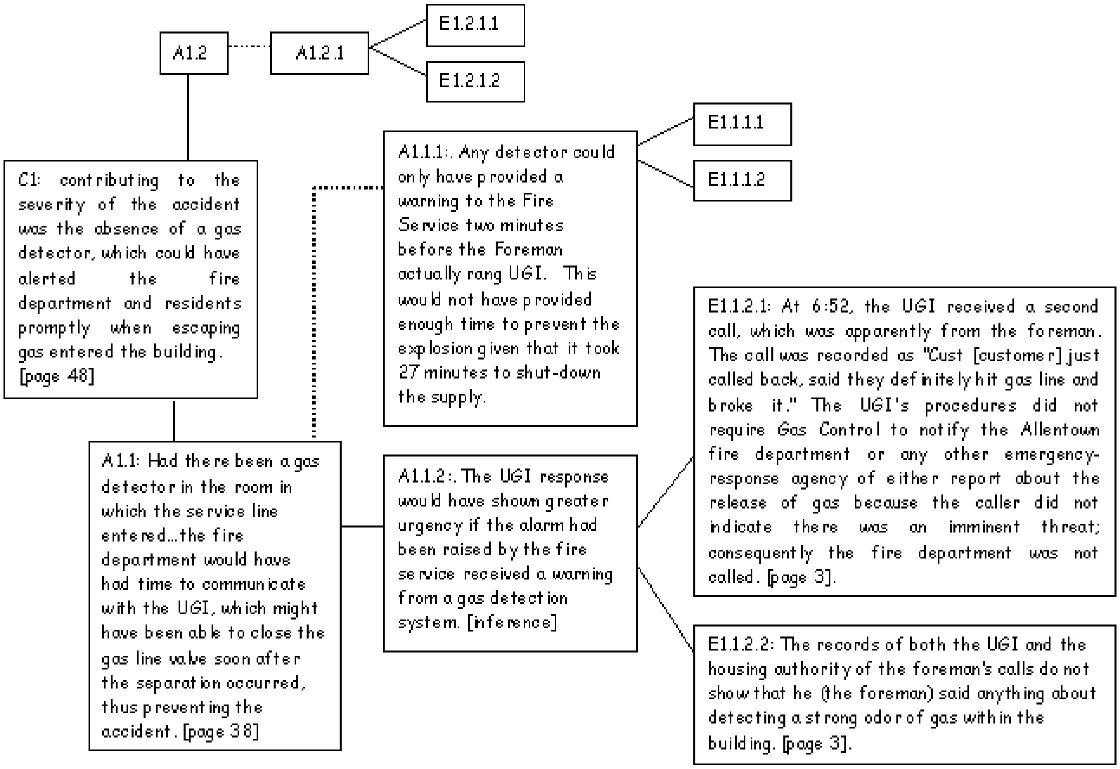

C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]

A1.1.1:. Any detector could only have provided a warning to the Fire Service two minutes before the Foreman actually rang UGI. This would not have provided enough time to prevent the explosion given that it took 27 minutes to shut-down the supply.

A1.1: Had there been a gas detector in the room in which the service line entered...the fire department would have had time to communicate with the UGI, which might have been able to close the gas line valve soon after the separation occurred, thus preventing the accident. [page 38]

A1.1.2:. The UGI response would have shown greater urgency if the alarm had been raised by the fire service received a warning from a gas detection system. [inference]

E1.1.2.1: At 6:52, the UGI received a second call, which was apparently from the foreman. The call was recorded as "Cust [customer] just called back, said they definitely hit gas line and broke it." The UGI's procedures did not require Gas Control to notify the Allentown fire department or any other emergency-response agency of either report about the release of gas because the caller did not indicate there was an imminent threat; consequently the fire department was not called. [page 3].

E1.1.2.2: The records of both the UGI and the housing authority of the foreman's calls do not show that he (the foreman) said anything about detecting a strong odor of gas within the building. [page 3].

Figure 9.27: Representing Counter Arguments in a CAE Diagram (3)

develop counter cases that raise questions about some of the investigators' findings. Elsewhere this technique has been used more directly to identify inconsistencies, errors and omissions in incident reports [412]. For instance, we have shown that investigators have placed the same individual in two different locations at the same time. The resulting CAE diagrams have much in common with other techniques for communication formal reasoning, such as tableaux or proof trees.

9.2 Requirements for Reconstructive Modelling

Previous sections have introduced a number of abstract notations that can be used to reconstruct the events that contribute to adverse occurrences. The intention has been to provide a broad overview of techniques that avoid some of the current limitations that affect the simulation environments introduced in Chapter 8. In particular, these more abstract notations can, typically, capture both catalytic failures but also the more latent and managerial failures that contribute to major incidents. There are, however, a number of problems that frustrate the application of these techniques to support the reconstruction of adverse occurrences. For instance, a considerable amount of training may be required before domain specialists and incident investigators can exploit the formal proof techniques that were introduced in the previous section. In contrast, temporal extensions to Fault Trees can initially be easier to understand. However, the lack of any formal semantic can lead to disagreement about the interpretation of these informal annotations. The following pages, therefore, address some of these limitations and derive requirements that investigators should consider when selecting an appropriate reconstruction technique.

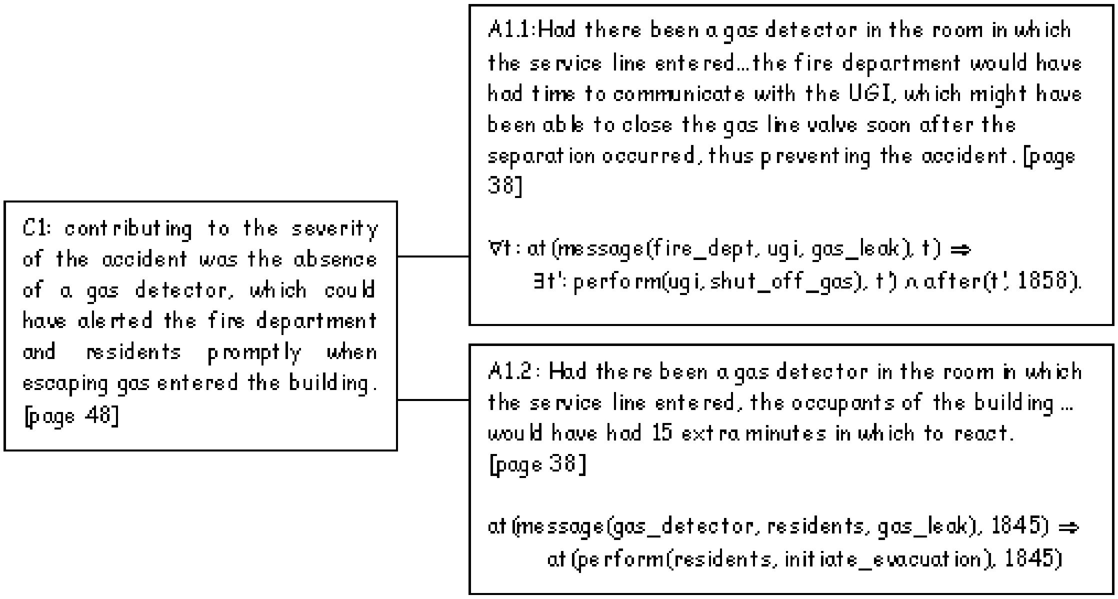
A1.1:Had there been a gas detector in the room in which the service line entered...the fire department would have had time to communicate with the UGI, which might have been able to close the gas line valve soon after the separation occurred, thus preventing the accident. [page 38]
∀t : at(message(fire_dept, ugi, gas_leak), t) ⇒
∃t': perform(ugi, shut_off_gas), t') ∧ after(t', 1858).
C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]
A1.2: Had there been a gas detector in the room in which the service line entered, the occupants of the building ... would have had 15 extra minutes in which to react. [page 38]
at(message(gas_detector, residents, gas_leak), 1845) ⇒
at(perform(residents, initiate_evacuation), 1845)

Figure 9.28: High-Level CAE Diagram Integrating Formal and Informal Material

9.2.1 Usability

Modelling notations must satisfy two different sets of requirements if they are to support incident reconstruction. The first centers on the usability of the technique; can investigators learn to apply the approach to quickly and accurately reconstruct the events leading to an incident? The second set of requirements focuses on expressiveness; does the notation enable designers to represent salient aspects of the incident?

Proportionate Effort and Ease of Learning

Different notations offer different degrees of support to various stages of the learning process. For instance, graphical notations may be easier for novices to understand than textual notations. Features such as a simple linear relationship between time and the position of annotations on a time-line can help people at the lower ends of the learning curve to focus upon key concepts rather than underlying mechanisms. Conversely the features of more advanced temporal logics, such as model based semantics and Kripke proof techniques, help more experienced analysts to exploit the full power of the language.

It is important to emphasise, however, that investigators will not invest the time necessary to gain additional expertise in complex modelling notations unless that are persuaded of the benefits. The rewards from using a notation must be perceived to be in proportion to the time taken to learn that notation [151]. This has significant consequences for some of the notations that have been introduced in this chapter. It has not been demonstrated that formal logics and semi-formal notations, included extended fault trees, offer significant benefits over less formal approaches, including graphical and textual time-lines. Unfortunately, this creates a paradox. More formal notations are rejected because they are not perceived to offer significant benefits. However, it is difficult to determine whether these approaches will offer significant benefits because they have not been widely adopted.

There have been a number of attempts to validate the potential benefits of semi-formal and formal notations both as tools for incident reconstruction and, more generally, to support the design of safety-critical systems. These studies yielded a number of interesting insights. For example, in one study we investigated whether engineers could learn to read and analyse complex reconstructions of safety-critical applications. The studies focussed on a number of different applications with complex

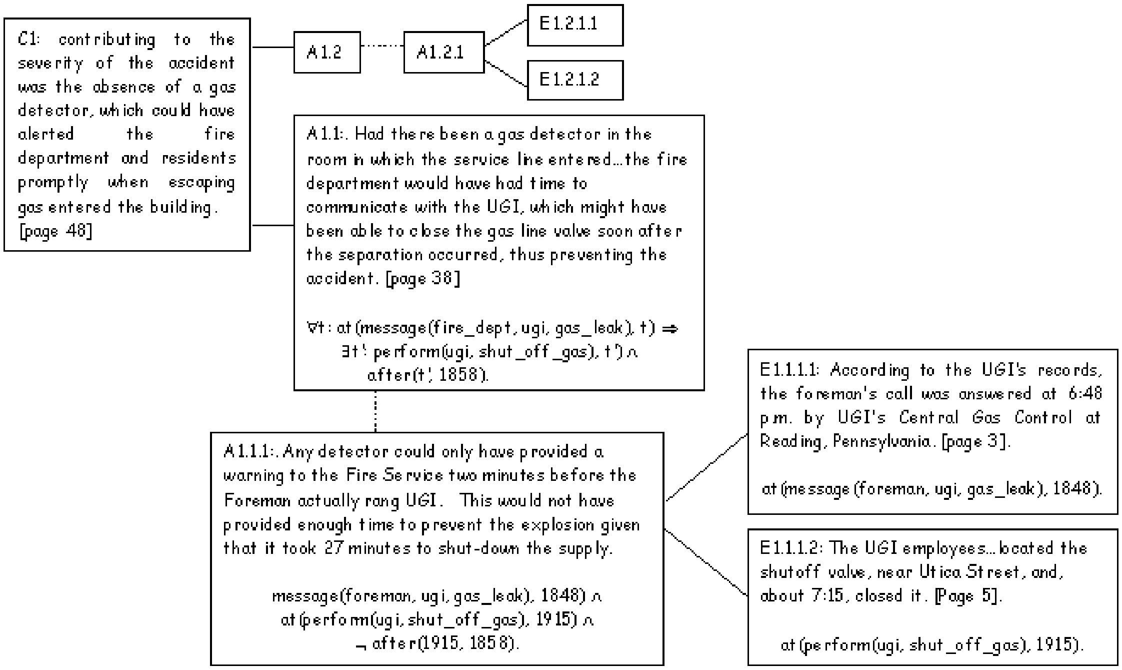

C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]

A1.1:. Had there been a gas detector in the room in which the service line entered...the fire department would have had time to communicate with the UGI, which might have been able to close the gas line valve soon after the separation occurred, thus preventing the accident. [page 38]

$$\forall t: at(message(fire\_dept, ugi, gas\_leak), t) \Rightarrow \exists t': perform(ugi, shut\_off\_gas), t') \wedge after(t', 1858).$$

A1.1.1:. Any detector could only have provided a warning to the Fire Service two minutes before the Foreman actually rang UGI. This would not have provided enough time to prevent the explosion given that it took 27 minutes to shut-down the supply.

$$message(foreman, ugi, gas\_leak), 1848) \wedge at(perform(ugi, shut\_off\_gas), 1915) \wedge \neg after(1915, 1858).$$

E1.1.1.1: According to the UGI's records, the foreman's call was answered at 6:48 p.m. by UGI's Central Gas Control at Reading, Pennsylvania. [page 3].

$$at(message(foreman, ugi, gas\_leak), 1848).$$

E1.1.1.2: The UGI employees...located the shutoff valve, near Utica Street, and, about 7:15, closed it. [Page 5].

$$at(perform(ugi, shut\_off\_gas), 1915).$$

Figure 9.29: Extended CAE Diagram Integrating Formal and Informal Material (1)

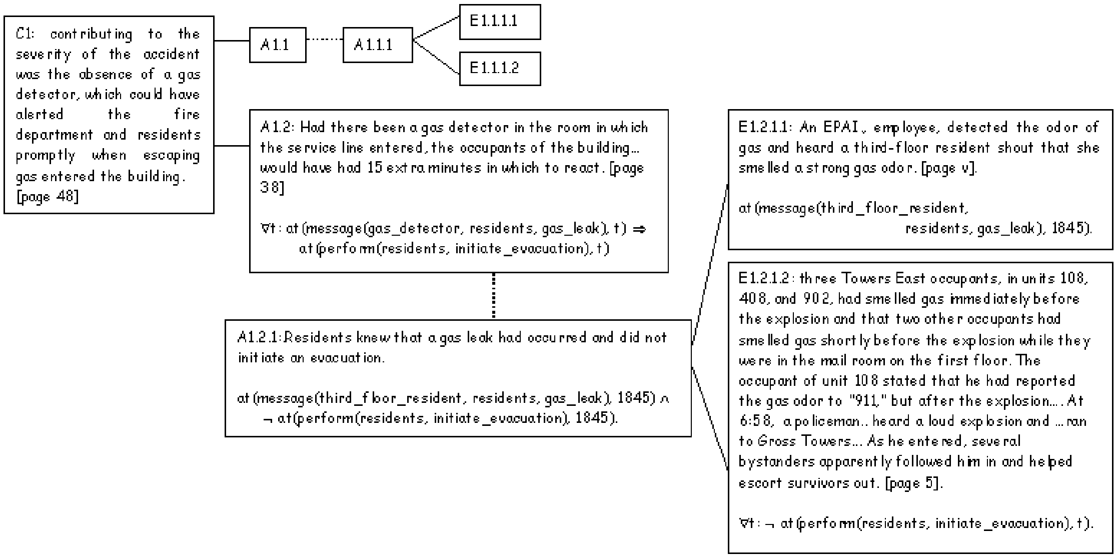

C1: contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building. [page 48]

A1.2: Had there been a gas detector in the room in which the service line entered, the occupants of the building... would have had 15 extra minutes in which to react. [page 38]

$$\forall t: at(message(gas\_detector, residents, gas\_leak), t) \Rightarrow at(perform(residents, initiate\_evacuation), t)$$

A1.2.1:Residents knew that a gas leak had occurred and did not initiate an evacuation.

$$at(message(third\_floor\_resident, residents, gas\_leak), 1845) \wedge \neg at(perform(residents, initiate\_evacuation), 1845).$$

E1.2.1.1: An EPAI., employee, detected the odor of gas and heard a third-floor resident shout that she smelled a strong gas odor. [page v].

$$at(message(third\_floor\_resident, residents, gas\_leak), 1845).$$

E1.2.1.2: three Towers East occupants, in units 108, 408, and 902, had smelled gas immediately before the explosion and that two other occupants had smelled gas shortly before the explosion while they were in the mail room on the first floor. The occupant of unit 108 stated that he had reported the gas odor to "911," but after the explosion.... At 6:58, a policeman.. heard a loud explosion and ... ran to Gross Towers... As he entered, several bystanders apparently followed him in and helped escort survivors out. [page 5].

$$\forall t: \neg at(perform(residents, initiate\_evacuation), t).$$

Figure 9.30: Extended CAE Diagram Integrating Formal and Informal Material (2)

failure models. These were modelled using temporal logic and a simplified Petri Net notation. This differed from the more convention notation introduced in this chapter because only one place was marked at any stage of a reconstruction. It was, therefore, very similar to state transition networks [403]. For instance, one study looked at the behaviour of a gas turbine controller. The participants were engineers stationed on rigs off the United States and Norwegian coasts. We faxed them example models and a number of associated questions. They were encouraged to take as long as they needed to answer the questions but to report the amount of time that they required to complete the questionnaires. They were expected to respond to two different types of question. The first tested their comprehension of the reconstruction. For instance, they were asked 'does the model describe any possible error condition after the application was loaded?' and 'was the application active after the error was acknowledged?'. The comprehension questions were counter-balanced so that subjects could not re-use their answers from the graphical reconstruction to answer questions about the logic model or vice versa. We also asked more qualitative questions about their impressions from using the formal and semi-formal notations. For instance, we asked them whether or not they would have preferred the reconstructions to have been expressed in natural language rather than the logic or the graphical notation.

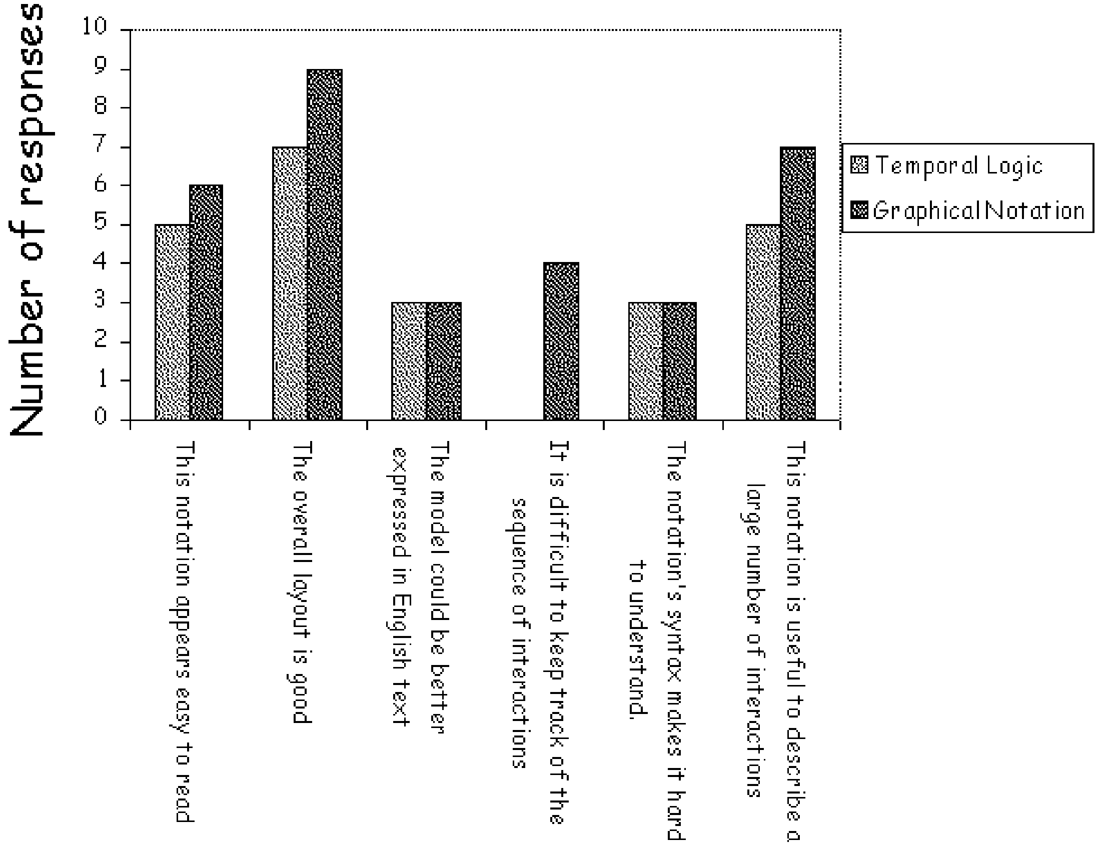

Figure 9.31: Subjective Responses to Modelling Notations.

The results confirmed many of our intuitions about the application of formal and semi-formal reconstruction techniques. For instance, the first set of fifteen US and Norwegian engineers only provided correct answers to 60 per cent of the comprehension questions using the graphical notation. The same group achieved a 55% success rate with the logic notation. Although these results seem disappointing, they were achieved without any formal training in the use of the notations. There were large deviations in individual scores. For instance, one engineer scored 100% in both conditions whilst another did not better than 30% correct. There were also some surprises. This group of engineers took an average of 8.2 seconds to answer the comprehension questions using the graphical notation and 8.7 seconds to answer using the logic-based reconstruction. Again, there were considerable deviations in individual performance. Figure 9.31 provides an overview of the responses to the modelling notations. Each individual had to tick a box stating that they agreed or disagreed with the statement. Each column, therefore, has a maximum value of 15. Perhaps the most surprising result here is that so few of the engineers believed that the model could be better expressed in natural language rather than either the graphical or logic based notation. This is interesting because it suggests that our limited sample of qualified engineers have a certain tolerance for the use of formal and semi-formal notations. Follow-up interviews revealed that similar techniques, for example fault trees, formed a common ingredient in their education and training.

There are a number of caveats that must be raised about such attempts to assess the usability of incident reconstructions. Previous chapters have argued that questionnaires and self-reporting techniques both raise a host of methodological questions about the reliability of the data that they yield. In particular, this initial study focussed on individual responses to a single set of tasks. It did not study the effectiveness of a reconstruction technique for the team-based tasks that typify incident investigations. Nor did it assess whether the long-term benefits of using either the graphical or logic-based technique were perceived to outweight any training overheads. In other words, it provided a single snap-shot of engineers' attitudes at a relatively early stage on the learning curve. It should also be stressed that our findings are not statistically significant, with limited exceptions [403]. Further studies are required to replicate these findings for other incident reconstructions and for greater numbers of potential users. The sample used in this study was relatively small. This was a consequence of our decision to use practicing engineers with similar skills and backgrounds. A number of practical reasons motivate the decision to restrict our sample in this way. Incident reconstructions must account for the technical causes of systems failure. It is, therefore, important that potential participants understand the potential causes of these systems failures. Otherwise, any results might stem from the participants ignorance about the application domain rather than from attributes of the reconstruction. However, this decision raised further issues. In particular, we could not obtain access to enough individuals with experience as incident investigators. Hence the exercise relied upon the participants' experience as design engineers attempting to diagnose potential problems that they had observed in a system rather than as incident investigators responding to reports from others within their organisation. Some of these caveats can be addressed by recruiting a larger group of participants. For example, a cohort of undergraduate students might have been used. However, the findings of such a study cannot easily be generalised to account for the attitudes of individuals who are likely to participate in incident investigations. Ultimately such studies probably require the financial backing and administrative support of regulatory authorities if they are to produce satisfactory results. We are, however, unaware of any field trials or studies that are specifically intended to validate potential techniques for incident reconstruction and modelling.

Visual Appeal

The previous section has argued that investigators must be persuaded of the practical benefits of reconstruction techniques if they are to invest time and money in learning to exploit them. The initial 'visual appeal' of a notation has a profound impact upon whether or not such investments will be made. For instance, logics are often rejected as being unnecessarily complex [427]. They lack the visual appeal of many graphical notations. However, this initial assessment can be very misleading. It can be difficult to maintain fault trees that extend to several hundred events. In contrast, mathematical abstraction techniques can be used to support the maintenance of large scale logic reconstructions [118]. It can be argued that the visual appeal of graphical notations must be weighed against the reasoning power of textual notations. This would, however, be too simplistic an analysis. For instance, there are strong text-based reasoning techniques associated with Petri Net reconstructions [678]. There are also well-established techniques for moving between these different representations. For example, Hura and Attwood demonstrate that the gates of a fault-tree can be represented by the places and transitions of a Petri Net [377]. Alternatively, the findings of formal proof techniques can be presented using semi-formal approaches that include the CAE diagrams and proof trees of previous sections.

There are additional costs associated with hybrid techniques that move between textual and graphical approaches or between formal and semi-formal notations. For example, it can be difficult to ensure that these multiple representations remain consistent during the course of an investigation. It is, therefore, again important to demonstrate the 'real-world' benefits of such hybrid techniques. We have conducted a number of studies to determine whether engineers can use semi-formal argumentation structures, similar to CAE diagrams, to address the usability problems that are often perceived to jeopardise the use of logic-based notations. For instance, a week-long trial was conducted with a group of software engineers from a range of industries. During this period the subjects were trained from 'scratch' to a level where they could both read and write logic-based models of complex, safety-critical systems. The first four days included an intensive course on discrete mathematics. On the fifth day, they were presented with a logic-based model of a control application for a chlorine recovery system. Elements of this model were then used to reconstruct the events leading to a previous incident involving this application. The engineers were asked a number of qualitative questions about the usability of the formalisation. The results of this are shown as Figure 9.32. As can be seen, our subjects found the model to be either impossible or hard to understand even after a week's intensive training.

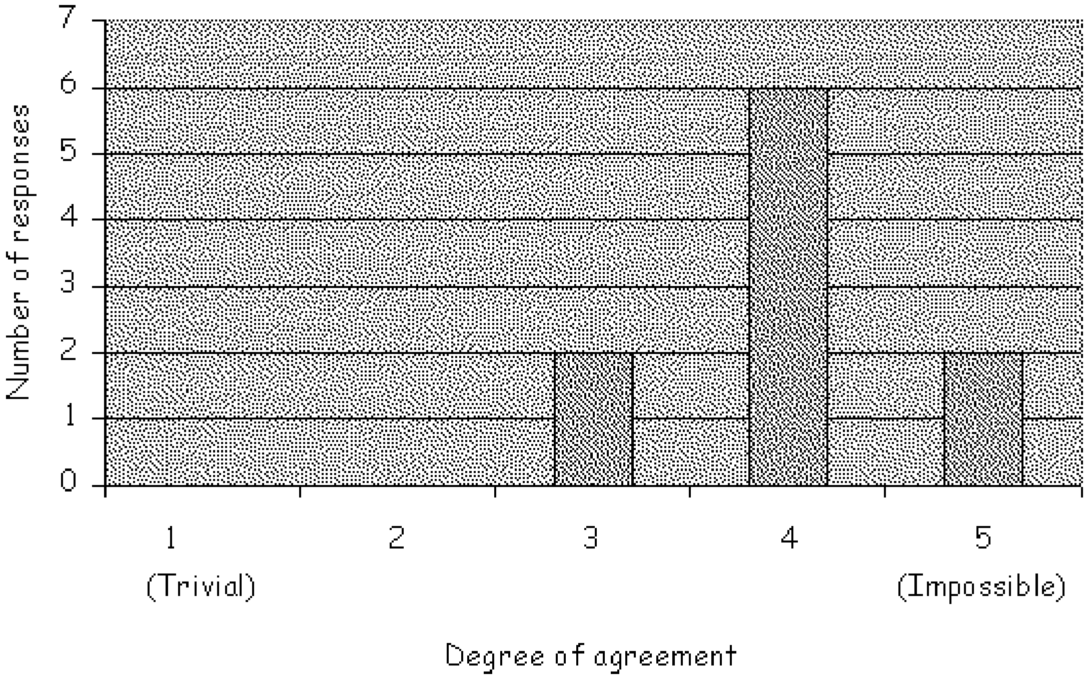

Figure 9.32: Subjective Responses to Logic-Based Reconstruction
How Easy did you find it to understand the logic-based model?

Such results are not particularly surprising. The application of logic is a skill based activity. The example used was of 'industrial strength'; it was based around the failure of a real system One week provides insufficient training to develop the expertise that is necessary to become confident in the use of formal reconstructions. Perhaps, more surprising are the qualitative responses for the semi-formal diagrams. After being asked to analyse the logic model, our subjects were shown a CAE-based diagram for another area of the chlorine recovery system. The ratings for this are shown in Figure 9.33. It should be noted that the logic-based reconstruction provided a detailed explanation of the events leading to an incident. In contrast, the graphical representation sketched the arguments for and against two competing explanations for a failure elsewhere in the recovery application.

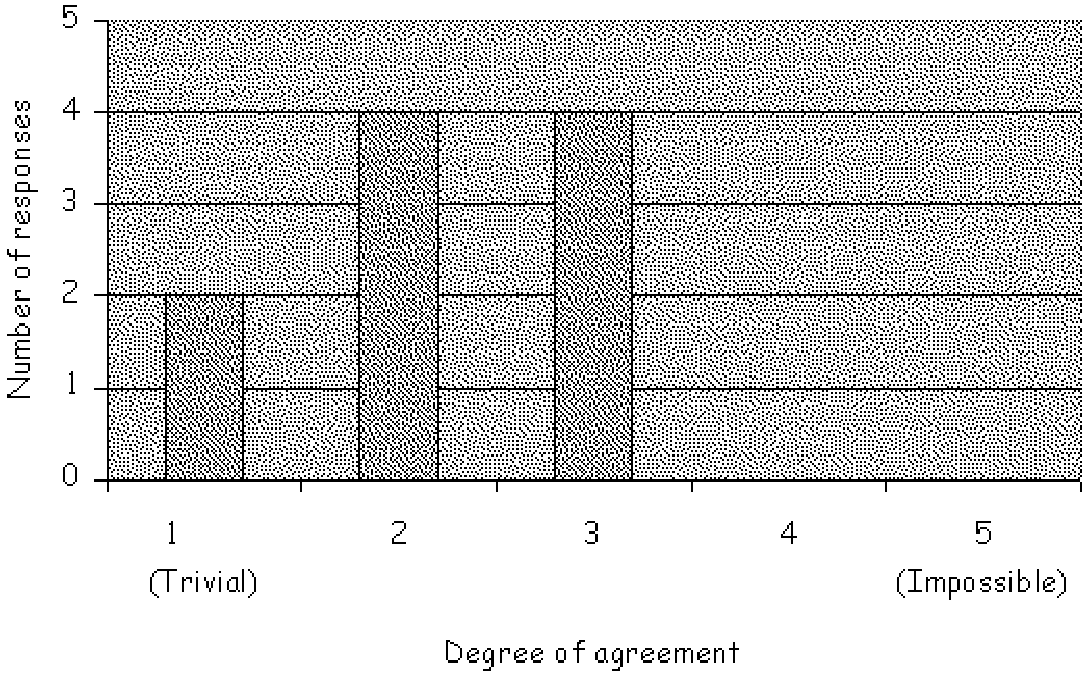

Figure 9.33: Qualitative Assessments Of CAE-Based Diagrams
How Easy Did You Find It to Understand the CAE Diagram?

In contrast to the formal specification, the subjects found it far easier to understand the graphical notation. It is important to emphasise, however, that no direct comparisons can be made between the attitude statements in Figures 9.32 and 9.33. Clearly, the information content is quite different. We then presented the participants with a more integrated reconstruction that that included logic clauses within a CAE diagram. The resulting diagram was similar to that presented in Figures 9.29 and 9.30. The participants' responses to this hybrid approach are shown in Figure 9.34. This provide some encouragement, especially considering the antipathy to logic-based reconstructions and that the participants had not any previous training in discrete mathematics.

As with the previous validation, these findings are suggestive rather than conclusive. we have not, to date, been able to guarantee the participation of a reasonable sample of trained incident investigators. As a result, we have been forced to rely upon the support of practicing engineers who have participated in incident investigations but who are not specifically trained in the investigatory techniques mentioned in previous chapters. There rae many reasons for this. One is that there are still relatively few trained investigators within even large-scale commercial organisations. They tend to be senior staff. In consequence, it can be difficult to secure their participation in such validation exercises. This study proved to be particularly difficult because it did not simply rely upon the one-off questionnaires that were described in the previous study. We had to train our subjects over a significant period of time; this involved a high degree of commitment from both the individuals concerned and from their companies. We are currently attempting to replicate our results with

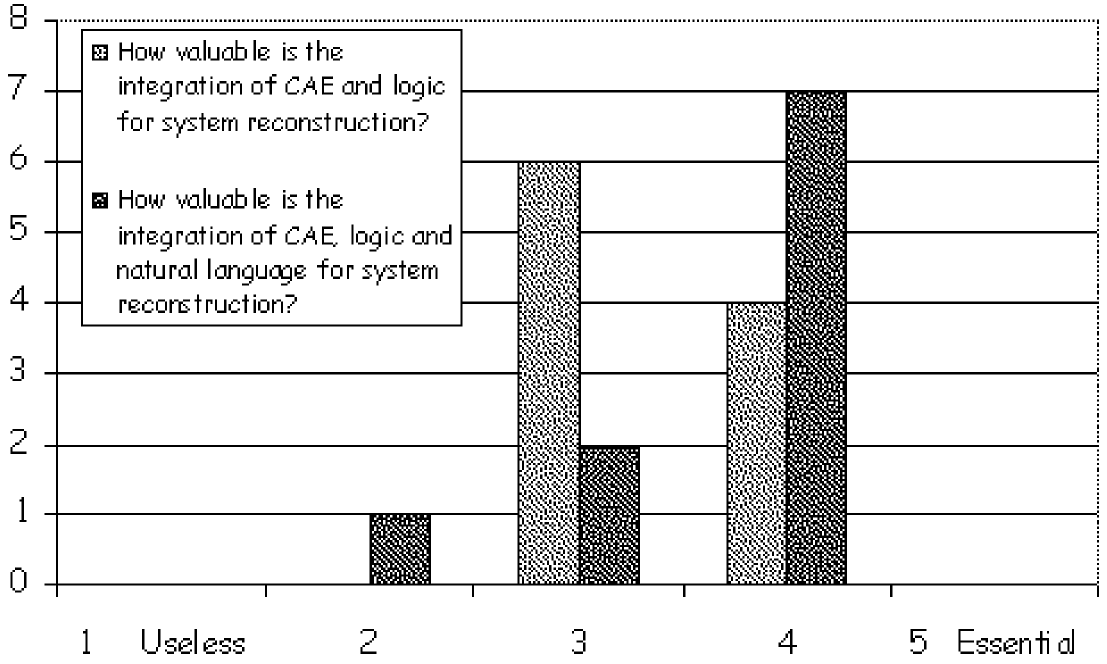

Figure 9.34: Qualitative Assessments of Hybrid Approach

larger groups of engineers and investigators. Again, however, it is difficult to foresee how many of the practical barriers will be resolved without greater regulatory commitment and support.

Tool Support

Previous sections have argued that semi-formal and formal notations provide investigators with means of focusing on critical properties of incidents and accidents. Irrelevant details can be stripped out to represent those events that contribute to an adverse occurrence. However, we have also demonstrated that these models can also become difficult to develop and maintain. For instance, there are significant overheads involved in constructing Petri Nets such as that shown in Figure 9.1.3. It can also be difficult to prove that the introduction of clauses, such as (9.2) and (9.3), does not contradict previous assertions about the course of an incident. It can be argued that this complexity is an inevitable consequence of our increasing desire to adopt a 'systems' approach to incident investigation. As we have seen, reconstructions must capture both the proximal and the distal causes of adverse occurrences. This inherent complexity helps to increase the importance of tool support during incident reconstruction. Tools can help in a number of ways. They can provide electronic support for the problems of constructing, navigating and typesetting complex graphical structures that might otherwise extend over many printed pages. Tool support can also implement syntactic checks to ensure that designers have constructed valid sentences from the lexical tokens in a formal language. They can partially automate reasoning about critical properties of incident reconstructions. As we have seen in the previous chapter, they can also be used to develop interactive simulations of adverse occurrences. Some tools enable these simulations to be directly derived from the abstract models that we have presented in this chapter.

The use of a formal or semi-formal notation does not guarantee the error-free development of an incident reconstruction. In Chapter 3 we defined a mistake to 'stem from a failure to select appropriate objectives irrespective of whether or not the actions taken to achieve those objectives are successful'. It is entirely possible that other analysts will conclude that investigators are mistaken in those aspects of an incident that they choose to reconstruct. We defined slips and lapses to 'result from some failure in the execution of a plan or well understood sequence of actions regardless of whether that plan was or was not appropriate'. By extension, it is also possible for investigators to develop a reconstruction that does not model an incident in the manner that they intended. For example, the structure of a Petri Net may make it impossible for places to be marked in the sequence that was intended by the investigator. Alternatively, a fault tree might have a minimal cut set that was not intended by the analyst and which could not have led to the incident given the available evidence.

At a higher level, it is possible for analysts to combine the tokens of a language to construct a model that has no meaningful interpretation. For instance, the places of a Petri Net must be connected to transitions. It is unclear what it would mean for one place to be connected directly to another place in such a graph. However, it can be difficult to avoid such errors when reconstructions can grow to include several hundred nodes or clauses. As a result, it is important to provide as much support as possible during the development of incident models. Type checking tools can ensure that relations hold between variables of the correct sort within the clauses of a logic model. Similar tools exist for the construction of both Petri Nets and Fault Trees. Without such support, it is difficult to conceive of large teams of designers constructing and maintaining detailed models of complex incidents. Computer-based tools can also conduct syntax checks. For instance, structure editors enable analysts to automatically insert syntactically correct components into a reconstruction. This raises a number of further usability issues. Some tools force analysts to *always* construct valid models. This can lead to considerable frustration. For example, it is frequently the case that investigators will have identified an important transition within an incident reconstruction. However, it may not be clear where it fits within the developing model. A tool that ensures continual correctness would force the analyst to link the transition to the rest of the network even if they did not feel confident about this placement. Incremental checking tools avoid this problem. They enable analysts to construct syntactically incorrect models. Places may initially be unconnected to any transitions and vice versa. However, these tools typically enable their users to periodically check the syntax of their structure once they feel confident that they have achieved a satisfactory placement of a node or that they have correctly constructed the axioms of their model. The meta-level issue is that not all tools provide equal degrees of support for incident reconstruction. Poorly designed tools may do little to address the usability problems that affect formal and semi-formal notations.

As mentioned, there are many different tools that can be recruited to support the reconstruction of complex incidents. The previous paragraph focussed on syntax editors and type checkers. However, other systems can be used to 'directly' develop prototype implementations from formal models [719, 720]. Chapter 8 included an example of a datalink air traffic control system that was simulated using this approach in Figure 8.15. This is important because formal and semi-formal notations can provide an extremely poor impression of the events leading to an accident. Interactive simulations can be shown to other analysts in order to validate the assumptions that are contained within accident models.

9.2.2 Expressiveness

A principle requirement for any incident reconstruction is that it should be capable of representing the diverse events that contribute to adverse occurrences. This creates problems because the temporal properties of control systems are very different from those of their operators. Similarly, the catalytic failures occur on a very different timescale to the period over which management and regulatory changes can be effected. For example, the NTSB investigators recorded the EPAI's foreman's recollections that:

> "The foreman said that he then went to his pickup truck and, using his cellular phone,2 called the gas company and the housing authority, telling them that he was excavating near the gas line and smelled gas. He stated that he next made three attempts to phone '911'. He said that each time he called, there was no answer. He said he then moved his truck to another spot in the parking lot in case the phone signal to his cellular

> phone was being blocked. He said that at the new location he again tried unsuccessfully to call '911'." [588].

This can be contrasted with the level of detail in the following observations about systems behaviour within the Cullen report into the Piper Alpha incident [193]. Here the focus is upon the observable behaviour of a gas detection system during the disaster:

> "It became apparent that only the larger leaks could give a flammable gas cloud containing the quantity of fuel evidently necessary to cause the observed explosion effects. Interest centred therefore particularly on series 42, which was the only test at a leak rate of 100 kg/min. In this test the low level alarms occurred first for C3 in 5 seconds, then for C2, C4 and C5 in 15, 20 and 25 seconds respectively..." (page 77).

The first quotation is based around an individual's recollections. The timings are vague and, in this case, difficult to substantiate. The second quotation provides clear and precise timings for alarms that have been validated by empirical studies on replicas of the system. These examples illustrate how the range of temporal properties that must be captured in any reconstruction is determined by the nature of the incident that is being considered. For example, the NTSB investigators did not consider it necessary to model the flow of gas within Gross Towers to the same level of detail as the enquiry team did for the Piper Alpha accident. However, the nature of the temporal properties being represented within any reconstruction is also determined by the evidence that is available from any primary or secondary investigation. Some timings can be grounded while other temporal information may be vague and imprecise. For instance, table 9.2 shows how the Foreman's recollections can be measured against the records of his cellular operator.

Time	In/Out/ Complete	Duration (secs.)	Connected (secs.)	Time Between Calls (secs)	Location
18:46:41	Out Complete	13	25	9	UGI Switchboard
18:47:15	Out Complete	87	95	5	UGI Emergency Number
6:48:55	Out Complete	70	82	15	Home, EPAI V.P.
18:50:32	Out Complete	39	47	173	UGI Emergency Number.
18:54:10	Out Complete	84	121	170	Housing Authority Answer Service
18:59:01	Out Incomplete	0	55	2	911 (Allentown)
19:00:02	Out Complete	162	175	3	Home, EPAI V.P.
19:03	In Complete	120	120	40	Not Recorded
19:05:40	Out		540		Private No.
19:14	In Complete	180	180		Not Recorded

Table 9.2: Cellular Phone Records for Allentown Foreman

The Beginning and the End

When does an incident actually begin? Previous sections have argued that this is a non-trivial question and it is worth reviewing the issue in the light of our case study. For instance, we have shown how the catalytic events centre around the operation of the backhoe and other heavy equipment during the removal of the soil. However, the incident could not have occurred if the pipeline had not been left relatively unsupported after the initial operation to remove the tank. Alternatively, the incident might have started when the EPAI foreman and crew were briefed for this particular operation or when their training missed necessary information about OSHA excavation requirements. At a more general level, this incident might have stemmed from the long-running discussions about Excess Flow Valves that were chronicled in Figure 9.1. The key point here is that the starting point for an incident is often a subjective decision that reflects the analyst's view of its causes. Incident modelling notations must, therefore, represent this subjective decision. It must be possible for readers to clearly identify the moment at which an analyst considers an incident to begin.

A related question is 'when does an incident end?'. As we have seen, many conventional risk analysis techniques stop with an undesired event. This is illustrated by the fault tree in Figure 9.8. As we have seen, however, incident reconstructions must also consider what happens after such an event. In particular, they must represent the way in which people and systems either exacerbate or mitigate the consequences of any failure. For example, the Police Officer played a key role in evacuating the survivors after the initial explosion. Similarly, the prompt response of the Fire Service and the medical agencies helped to ensure that the injured were swiftly evacuated from the scene of the incident. These actions did not *cause* the accident but they did contributed to the saving of lives. They reduced the consequences of the failure itself.

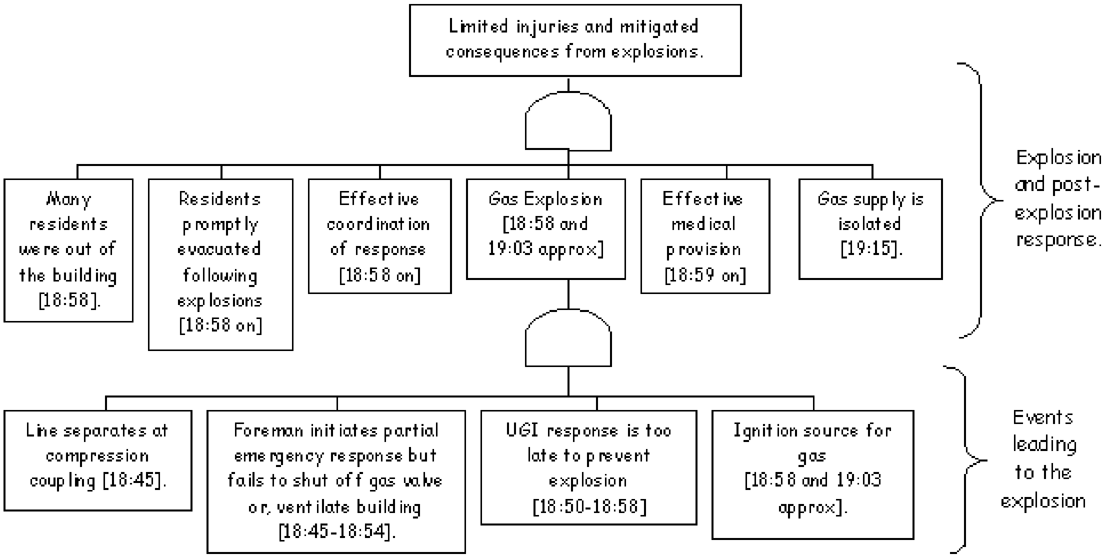
Limited injuries and mitigated consequences from explosions.
Many residents were out of the building [18:58].
Residents promptly evacuated following explosions [18:58 on]
Effective coordination of response [18:58 on]
Gas Explosion [18:58 and 19:03 approx]
Effective medical provision [18:59 on]
Gas supply is isolated [19:15].
Explosion and post-explosion response.
Line separates at compression coupling [18:45].
Foreman initiates partial emergency response but fails to shut off gas valve or, ventilate building [18:45-18:54].
UGI response is too late to prevent explosion [18:50-18:58]
Ignition source for gas [18:58 and 19:03 approx].
Events leading to the explosion

Figure 9.35: Allentown Fault Tree Showing Pre- and Post-Incident Events

Figure 9.35 uses elements of the Fault Trees that were constructed in previous sections of this chapter to show how modelling notations can be used to reconstruct events leading to, and stemming from, an adverse occurrence. The analyst's view of the start and finish of the accident are explicitly bounded by the extent of the tree. In this case we have not expanded the events that, for instance, contributed to Foreman's response. If investigators considered that such details fell within the scope of any analysis then they could be introduced as shown in Figure 9.10. Nor does it expand on the events that contributed to the effective coordination of the Emergency Services' response from 18:58 on, illustrated in Figure 9.13. These details can, of course, be introduced to explicitly indicate that they fall within the scope of the investigation. The development of the incident fault tree, therefore, encourages analysts to represent the extent of their enquiries. This can help to avoid the implicit decisions and misunderstandings that may threaten any subsequent causal analysis.

Figure 9.35 illustrates the strengths and the weaknesses of fault trees as a reconstruction notation. Te scope or extend of incident is explicitly represented. However, the lines between nodes represent a mixture of causal, temporal and logical relationships. This overloading provides considerable expressive power. It can also be misleading. For instance, previous sections have argued that incident fault trees can be formatted to preserve a left to right temporal ordering. Events and gates that occur during the early stages of an incident should be drawn to the left of components that occur later on. However, this convention does not form any part of the syntax or semantics of the fault tree notation. We have also shown the problems that arise when attempting to satisfy such a requirement. Events at one level in a tree can occur after or before events at another level. The best example of this is where some event in the aftermath of an incident is influenced by another, organisational or managerial, event that occurred long before the incident took place. In this case, the organisational event that contributed to the response would be shown higher-up the tree because its relevance is not to the pre-incident events but to the consequences of that failure.

Concurrency

Figure 9.36 illustrates the structure of many incident reports. Each chapter presents a chronology of events from a different perspective. A synopsis or overview chapter is followed by an analysis of any systems failure. The systems analysis is followed by an investigation of operational and management issues. This, in turn, is followed by an interpretation of any emergency response and so on. As a result, if a reader wants to build up a coherent view of all of the events in an incident at a particular point in time then they are forced to cross-reference many different sections of the report. For example, the events occurring at times T1 and T2 are described in each of the chapters represented in Figure 9.36. These problems also affect investigators during the stages of reconstruction and analysis that precede the drafting of an incident report. They must piece together information about the many different aspects of complex systems failures. This implies that there must be some means of representing and reasoning about concurrent interaction between the simultaneous failures that contribute to many incidents and accidents.

As we have seen, there are a range of graphical and textual notations that can be used to address these concerns. They provide explicit means of representing the concurrent events that occur in different areas of a system. They can also be used to represent the way in system failures and human error combine, at critical moments, to create the circumstances for an accident. To illustrate the importance of this, consider the following excerpts provided by the NTSB investigators into the Allentown incident:

> "When an Allentown fire inspector was inspecting the EPAI's work, he saw the excavation's west sidewall slide into the excavation, exposing the gas line, which was 3 to 4 feet west of the tank. The collapsed sidewall removed the soil support from about 30 feet of gas line causing it to sag." ([588], page 10).
>
> "Neither the EPAI employees nor the fire inspectors notified the UGI that the service line was unsupported and damaged. Later on May 23, the EPAI crew placed a cable sling around the tank and attached it to a chain that was attached to the backhoe. When the crew tried to lift the tank, the chain broke. Those who witnessed the event, including the second fire inspector, stated that they did not believe the tank struck the gas line". ([588], page 11),
>
> "Because the citys fire inspectors saw on May 23 that the service line was unsupported, they could have prevented the accident. They showed proper concern about the safety of the line, especially after a piece of asphalt pavement fell on it and deformed it. However, not having been instructed to do otherwise, both inspectors relied on the EPAI foremen's assessment that the line was safe". ([588], page 36).

These quotations illustrate how it can often be difficult for readers to form a coherent model of the events that are identified during incident investigations. for instance, these different accounts

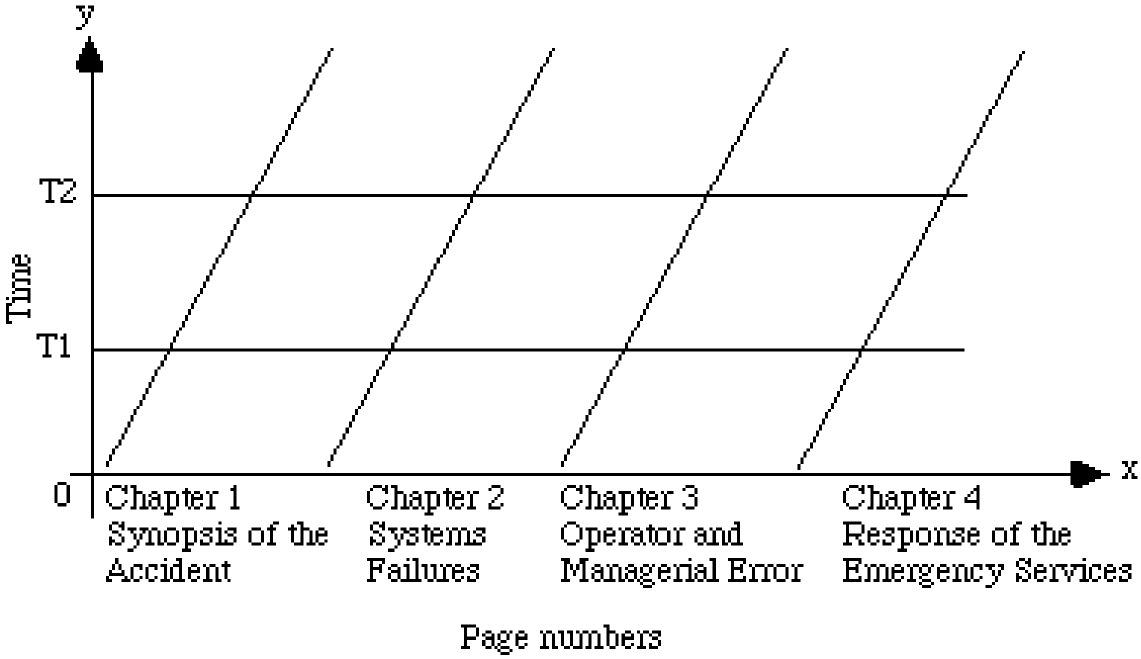

Figure 9.36: Cross-Referencing Problems in Incident Reports

do not state the order in which the wall collapsed, the chain broke or the asphalt struck the gas pipe. This ordering has to be inferred from evidence presented elsewhere in the report. Similarly, it can be difficult to determine how these different events affected the different people who were involved in the incident. Figure 9.37 builds on the Petri Nets that were introduced in previous sections to reconstruct a more coherent model of some of these events. As can be seen, the marking in this diagram denotes that the foreman is initially happy with the safety of the line, in spite of the inspector's concerns, and that asphalt is being lifted over the gas supply. The diagram, therefore, simultaneously captures human factors observations, derived from eye witness statements, together with information about the observable sequence of events leading to the incident.

There are a number of limitations with the previous diagram. There is little direct evidence to show that the asphalt strike triggered the Foreman's decision to support the pipe although this implied by the NTSB investigators. More significantly, however, Figure 9.37 only captures the relative timings of various events. The excavation slip occurred before the inspector questioned the Foreman about the safety of the gas line. The asphalt was being moved across the gas line before it was deflected and so on. What the previous diagram does not represent is the real-time at which these different events occurred. This is a significant limitation. For instance, there might have been seconds, minutes or even hours between the deflection of the pipeline and the Foreman's decision to reconsider the safety of their system.

Figure 9.37 illustrates the strengths and weaknesses of Petri Nets for incident reconstruction. This diagram shows how the notation can be used to generalise beyond the specific circumstances of a particular incident. Previous sections have argued that the temporal characteristics of previous incidents are unlikely to be exactly replicated in future failures. For example, there was a considerable delay between the failure of the excavation wall and the physical damage that separated the exposed pipeline to Gross Towers. In future, however, there might only be a matter of seconds between the excavation failure and direct physical damage to an exposed gas supply. The Petri Net illustrated in Figure 9.37 clearly avoids any commitment to such absolute timings that might not be replicated in future incidents. This ambiguity is, however, a significant weakness if investigators are concerned to accurately represent key properties of this particular incident. For instance, previous sections shown

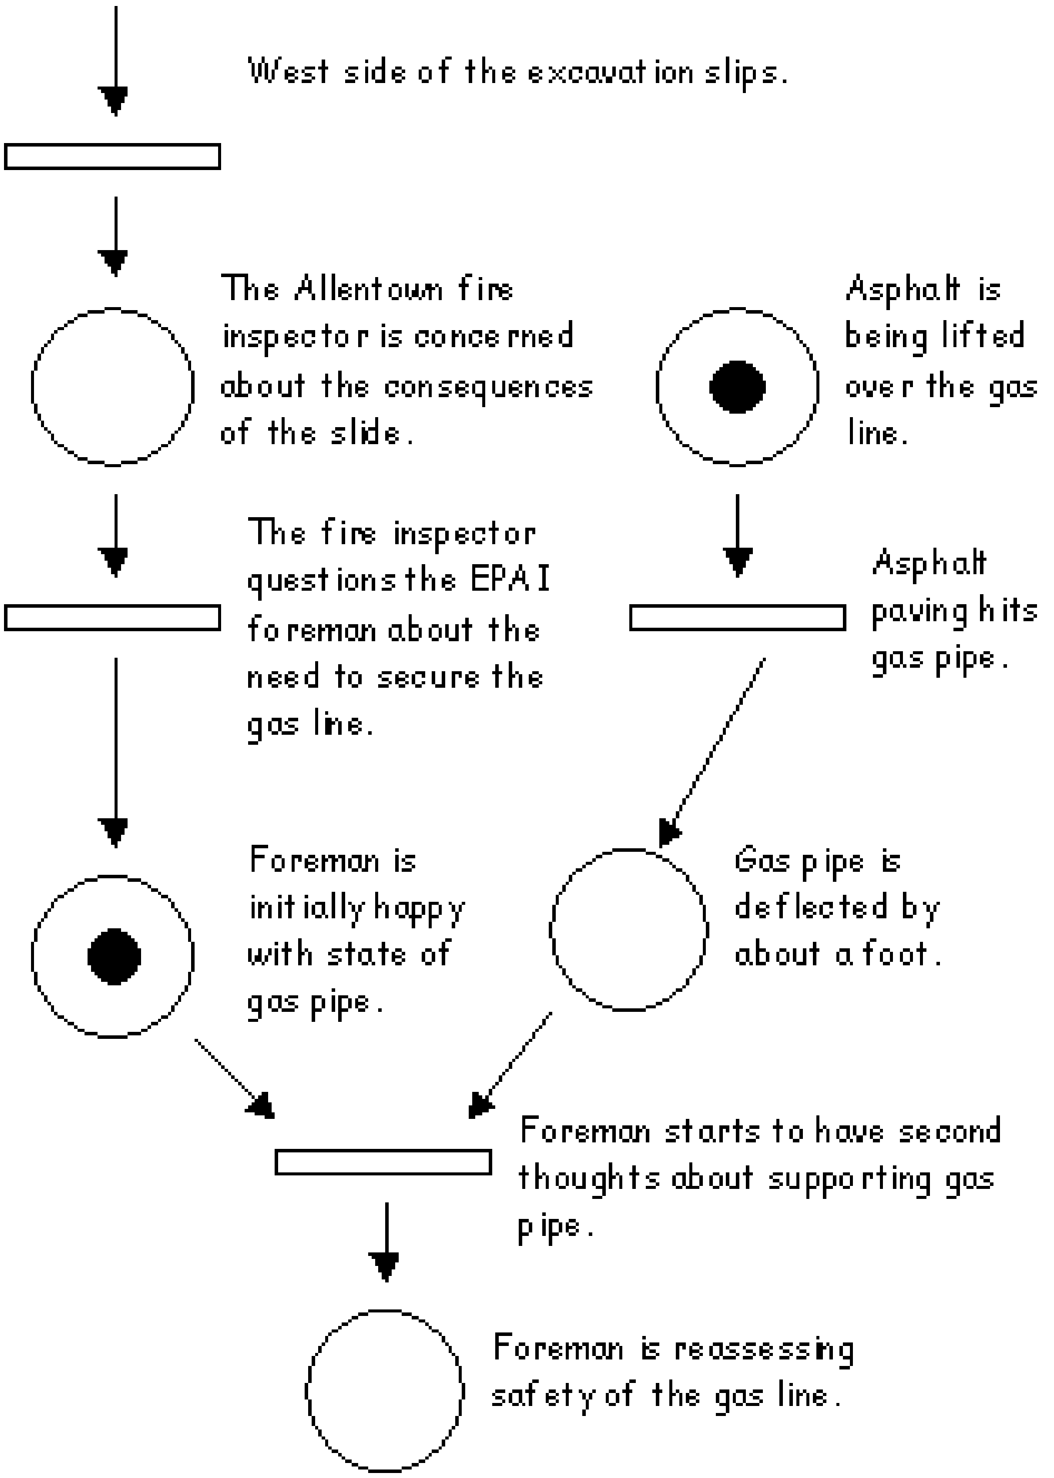
West side of the excavation slips.
The Allentown fire inspector is concerned about the consequences of the slide.
Asphalt is being lifted over the gas line.
The fire inspector questions the EPAI foreman about the need to secure the gas line.
Asphalt paving hits gas pipe.
Foreman is initially happy with state of gas pipe.
Gas pipe is deflected by about a foot.
Foreman starts to have second thoughts about supporting gas pipe.
Foreman is reassessing safety of the gas line.

Figure 9.37: Using a Petri Net to Build a Coherent Model of Concurrent Events

how investigators can temporal logics to examine the real-time characteristics of this incident. In particular, we have demonstrated that warnings from an automated gas detection system might not have prevented an explosion. Such an analysis cannot easily be performed using the relative sequences provided by the Petri Net in Figure 9.37.

Lack of Evidence

The previous section has made the case that incident modelling notations must be capable of representing real and interval time properties of adverse events. It is important, however, to emphasise that this must not force analysts into undue commitment when the exact timing for an event is unknown. For example, the NTSB investigators concluded that:

> "...the backhoe probably struck the line when being operated across it; the foreman's reports to both the UGI and the housing authority indicated that the pipe had been struck during recent excavation activities. Although the foreman denied after the accident that the backhoe had struck the line, the coating of the pipe showed evidence of mechanical damage, as did the pipe steel at one location. Also, the foreman's calls both to the housing authority and to the UGI show that at the time he believed his crew had hit the gas line while excavating." [588]

The use of terms such as 'probably' re-iterate the point that uncertainty often remains within models and reconstruction of safety-critical incidents. This uncertainty has many causes. For example, it may not be possible to obtain direct evidence to support the investigators' hypotheses. Alternatively, physical evidence can be contradicted by eye-witness testimony. In this case, the physical evidence of damage to the pipeline is contradicted by the foreman's recollections. Such contradictions can occur when witnesses do not observe key events during an incident. They can also result from the cognitive effects of stress, anxiety and guilt that have been discussed in previous chapters. This uncertainty can take many forms. For instance, the previous quotation centres on whether or not the backhoe struck the gas line when it was being operated across it. Even if we assume that the physical evidence does indicate that such damage was incurred then we cannot be certain of exactly when this happened. In consequence, even with sophisticated logging techniques it may not be possible to associate particular events with particular moment in time. Some notations provide more support for the representation of this lack of evidence than others. For example, time-lines may be extended with informal annotations as shown in Figure 9.38.

The annotations below the time-line are used to indicate the position of events whose time is known, either through corroborated eye witness statements or through external monitoring of the event. In contrast, the annotations above the line are used to indicate imprecise timings or events for which there is contradictory evidence. The horizontal parentheses under the label Backhoe probably strikes the gas line while being operated over it is used to indicate that the event occurred one or more times between 13.30-18.40. We do not know exactly when this occurred during this interval. Such annotations do not form part of the conventional time-line notation. This is important because analysts would have to learn to exploit a number of further extensions if such an approach were to represent the differing forms of temporal uncertainty that arise during many investigations. These can be summarised as follows:

- *a certain event with uncertain timing.* The event is known to have taken place but there is no clear evidence for when it occurred;
- *an uncertain event with uncertain timing.* It is not clear whether this event actually occurred or, if it did, when it actually took place. In some respects, this is the pathological case for incident reconstruction;
- *a certain event with certain timing.* This is the ideal case. There is clear evidence that an event occurred and there is evidence for when it took place.
- *an uncertain event with certain timing.* It is unclear whether the event actually occurred but, if it did, there is evidence for when it must have taken place.

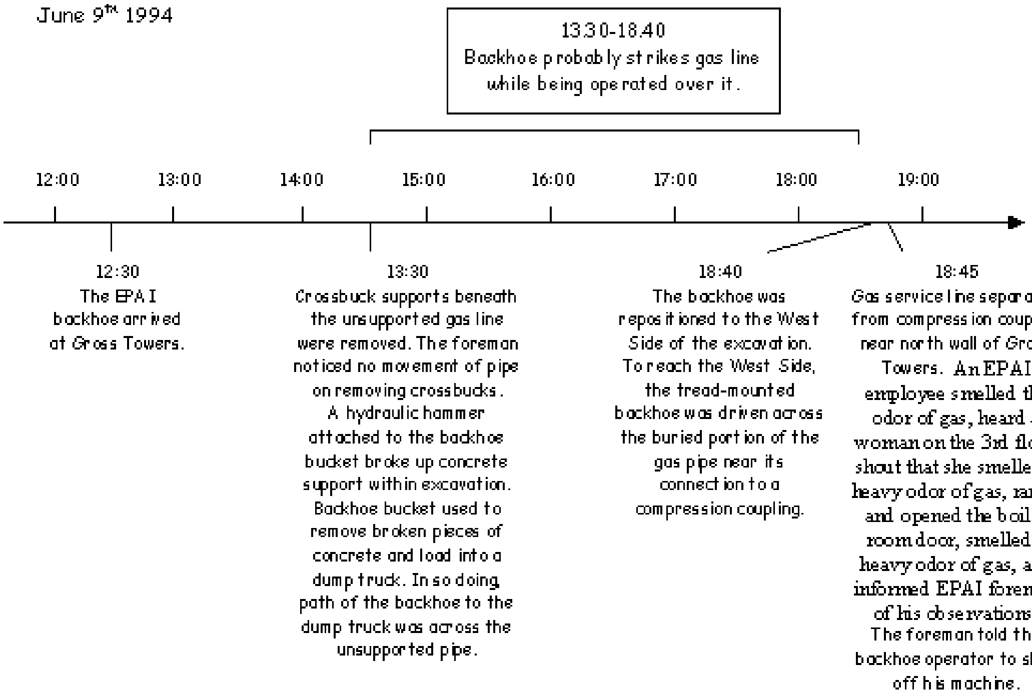
June 9th 1994
13:30-18:40 Backhoe probably strikes gas line while being operated over it.
12:00
13:00
14:00
15:00
16:00
17:00
18:00
19:00
12:30 The EPAI backhoe arrived at Gross Towers.
13:30 Crossbuck supports beneath the unsupported gas line were removed. The foreman noticed no movement of pipe on removing crossbucks. A hydraulic hammer attached to the backhoe bucket broke up concrete support within excavation. Backhoe bucket used to remove broken pieces of concrete and load into a dump truck. In so doing, path of the backhoe to the dump truck was across the unsupported pipe.
18:40 The backhoe was repositioned to the West Side of the excavation. To reach the West Side, the tread-mounted backhoe was driven across the buried portion of the gas pipe near its connection to a compression coupling.
18:45 Gas service line separated from compression coupling near north wall of Gross Towers. An EPAI employee smelled the odor of gas, heard a woman on the 3rd floor shout that she smelled a heavy odor of gas, ran to and opened the boiler room door, smelled a heavy odor of gas, and informed EPAI foreman of his observations. The foreman told the backhoe operator to shut off his machine.

Figure 9.38: Lack of Evidence, Imprecise Timings and Time-lines

Even this list is a simplification. For instance, investigators may have evidence that an event did occur and that it happened at a particular moment during an incident. However, there may not be any evidence about the duration of an event or if it occurred more than once. There are further complexities. For instance, it is important to distinguish between instantaneous events and more gradual changes that influence the underlying state of any system. There is an important distinction between this sort of information and that shown above the time-line in Figure 9.38. In the former case the event is instantaneous but it's timing is not known, in the latter case the property is continuous and its duration is well known. This distinction could be supported by introducing further annotations within the time-line notation.

Figure 9.39 illustrates the way in which additional syntactic features must be introduced to represent gradual changes in the underlying state of the system. In this instance, a different form of horizontal parentheses denote a continuous change over an interval rather than a discrete event at a particular point in the time-line. This diagram also illustrates the use of previous annotations to denote imprecise information. The text above the time-line is used to represent the lack of information about when exactly the foreman ordered his crew to trace the line back towards Utica Street. It is important to emphasise that the degree of uncertainty that is represented in diagrams such as Figure 9.39 will change over time. There is a strong motivation for investigators to resolve ambiguity as more evidence becomes available. Techniques, such as time-lines, that can be used to represent an event without commitment to whether it occurred or when it occurred are, therefore, more appropriate to the early stages of reconstruction. Other techniques, including computer-based simulation, that force greater commitment to particular timings are used more often in the later stages of an investigation.

Figure 9.39 shows how investigators must extend the basic time-line notation if they are to distinguish between different forms of uncertainty or between discrete events and continuous change. This illustrates the flexibility of this informal notation. The absence of strong syntactic rules enables designers to introduce novel features without worrying about whether or not the resulting diagrams

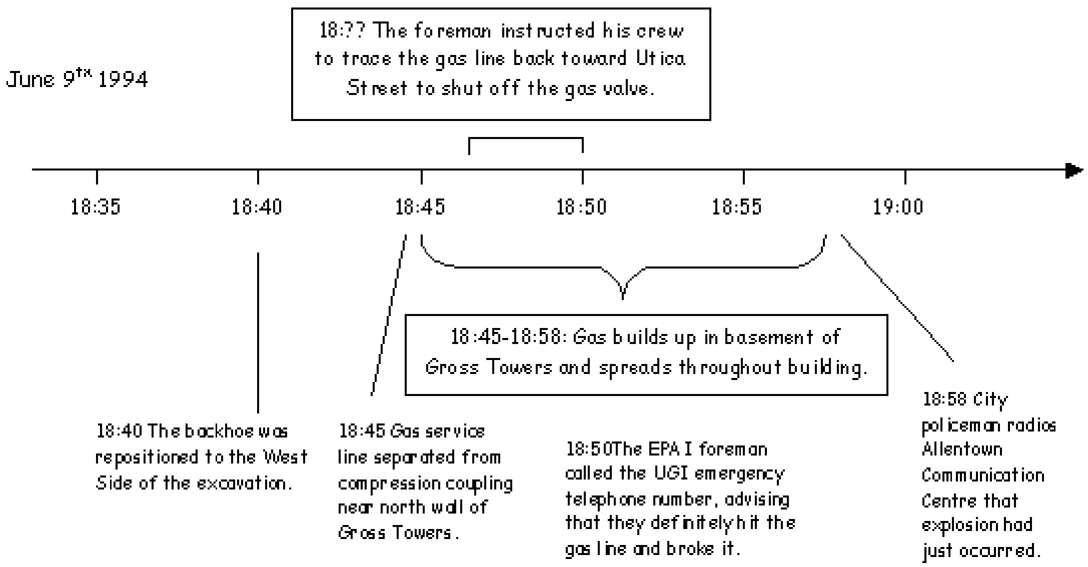
June 9th 1994
18:?? The foreman instructed his crew to trace the gas line back toward Utica Street to shut off the gas valve.
18:35
18:40
18:45
18:50
18:55
19:00
18:45-18:58: Gas builds up in basement of Gross Towers and spreads throughout building.
18:40 The backhoe was repositioned to the West Side of the excavation.
18:45 Gas service line separated from compression coupling near north wall of Gross Towers.
18:50The EPA I foreman called the UGI emergency telephone number, advising that they definitelyhit the gas line and broke it.
18:58 City policeman radios Allentown Communication Centre that explosion had just ocurred.

Figure 9.39: Continuous changes and Time-lines

will represent 'valid' or well formed time-lines. However, this freedom also results in a proliferation of ad hoc annotations within different investigation teams. During the prepartion of this book, I witnessed different investigators use the same annotation to represent different types of temporal properties. One used an asterisk to represent an uncertain event with a known time whilst their colleague used it to represent multiple occurrences of a known event at a known time. As a result, other members of the team had to recognise who had drawn any particular asterisk in order to know what it meant!

Inconsistencies

It is a frequent observation in incident reports that the evidence of one witness does not agree with that of another. Most often, these disagreements focus upon the sequence and timing of critical events. Alternatively, as we have seen in the previous section, they may disagree about whether or not those events ever took place at all. The following citation provides a further example of such contradictions. The foreman stated that he and other crewmembers supported the pipeline before they left the site. In contrast, housing authority employees testified that the line was unsupported:

> "The tank was successfully removed from the excavation, and samples of soil were taken adjacent to the tank's concrete support, which remained in the excavation. The soil was to be tested to determine whether fuel had leaked from the tank and contaminated the surrounding soil. The EPAI foreman stated that before he and the other crewmembers left the site, they tried to support the pipe with saw horses, surrounded the excavation with orange plastic barrier fencing, put plastic sheeting over the excavation slopes, including the soil that lay beneath the pipe, and removed the equipment from the site. They left the excavation open to await the result of the tests. Housing authority employees who frequently passed the excavation between May 23 and June 9 stated they observed that the exposed pipe was not supported." [588]

Analysts must consider the different scenarios that are created by such uncertainty. The following Petri Nets illustrate this point. The diagram on the left of Figure 9.40 presents an extract from the Petri Net previously introduced in Figure 9.1.3. This represents the view that the saw horses were left providing partial support for the exposed gas line after the excavation team left the site. In contrast, the Petri net on the right represents an alternative version of events based on the House

Association employees' testimony. This extends the previous networks by hypothesising that the unstable soil and adverse weather conditions contributed to the collapse of the supports that had previously been placed under the gas line. The main conflict arises between the Housing Association employees' observations and the testimony of the Foreman and his crew, confirmed by the two Inspectors.

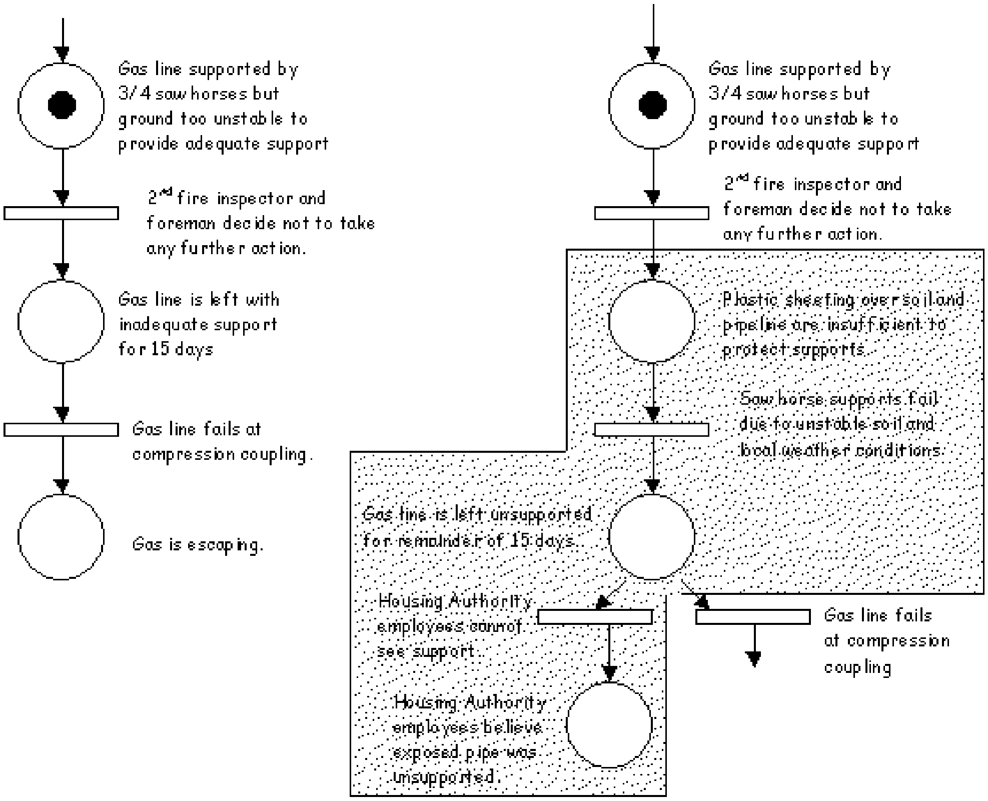
Gas line supported by 3/4 saw horses but ground too unstable to provide adequate support
2nd fire inspector and foreman decide not to take any further action.
Gas line is left with inadequate support for 15 days
Gas line fails at compression coupling.
Gas is escaping.
Gas line supported by 3/4 saw horses but ground too unstable to provide adequate support
2nd fire inspector and foreman decide not to take any further action.
Plastic sheeting over soil and pipeline are insufficient to protect supports.
Saw horse supports fail due to unstable soil and local weather conditions.
Gas line is left unsupported for remainder of 15 days.
Housing Authority employees cannot see support.
Housing Authority employees believe exposed pipe was unsupported.
Gas line fails at compression coupling

Figure 9.40: Using Petri Nets to Represent Different Versions of Events

Petri Nets have not previously been used to represent and reason about such inconsistency. This approach does not, however, provide an ideal solution. As we have seen, these networks can become extremely complex even for relatively simple behaviours. The problems associated with constructing and maintaining these diagrams can be exacerbated if they are used to represent multiple, alternative accounts of the same failures. Analysts must manually inspect the different networks in order to identify the differences that exist between these individual accounts. Figure 9.40 provides partial support by shading the area of the network to denote potential disagreement over the course of events. This is not, however, a general solution. For example, subtle differences of interpretation about the initial causes of an incident might have consequences that extend throughout any model or reconstruction. As a result, almost every node within a network might be shaded [406]. A more pragmatic solution is to find evidence that can be used to resolve any apparent contradictions.

Figure 9.41 shows how analysts can use the Petri Net notation to construct a third version of event that resolves the previous inconsistency. The Petri Net on the left shows that from certain positions around the excavation, the pipe might have appeared to be unsupported even though the saw horses were still in place. This can be compared with the Petri Net on the right of Figure 9.41. This network was introduced in Figure 9.40; supports failed at some point after the excavation team left the site. The key point is that the explicit reconstruction of an incident encourages investigators to identify and resolve potential inconsistencies. Additional evidence must be sought to determine which hypothesis is correct. Where there is contradictory evidence, the skill and

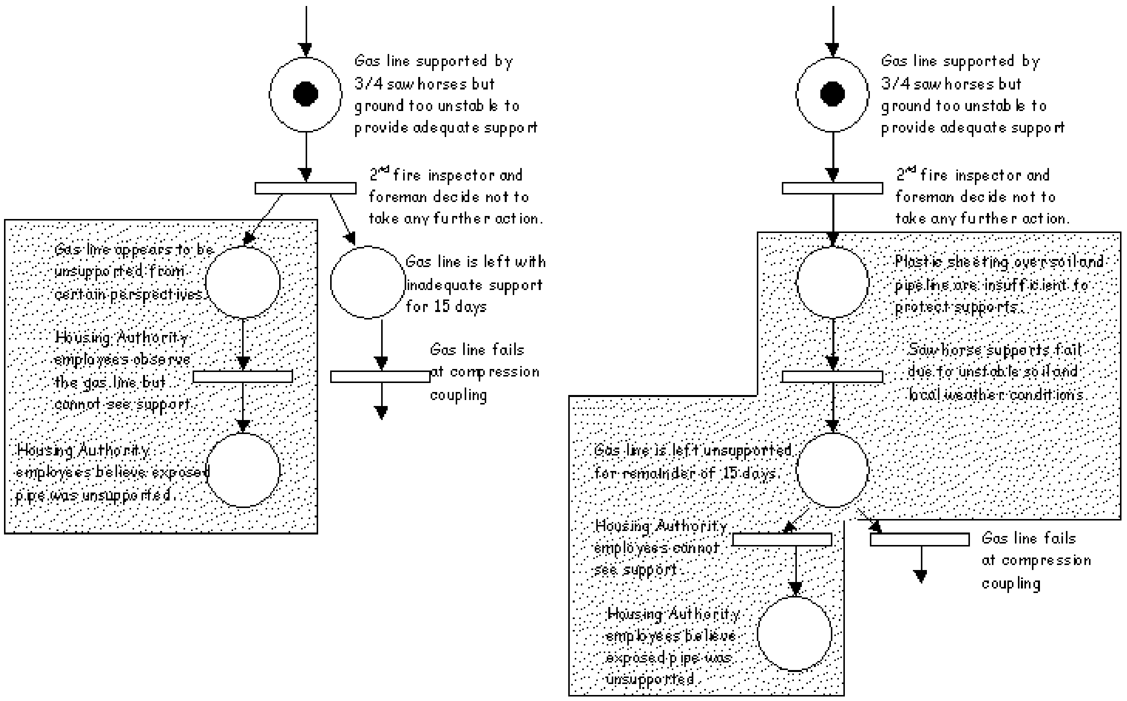
Gas line supported by 3/4 saw horses but ground too unstable to provide adequate support
2nd fire inspector and foreman decide not to take any further action.
Gas line appears to be unsupported from certain perspectives.
Housing Authority employees observe the gas line but cannot see support.
Housing Authority employees believe exposed pipe was unsupported.
Gas line is left with inadequate support for 15 days
Gas line fails at compression coupling
Gas line supported by 3/4 saw horses but ground too unstable to provide adequate support
2nd fire inspector and foreman decide not to take any further action.
Plastic sheeting over soil and pipeline are insufficient to protect supports.
Saw horse supports fail due to unstable soil and local weather conditions.
Gas line is left unsupported for remainder of 15 days.
Housing Authority employees cannot see support.
Housing Authority employees believe exposed pipe was unsupported.
Gas line fails at compression coupling

Figure 9.41: Annotating Petri Nets to Resolve Apparent Contradictions

judgement of the investigator must identify a 'probable' version of events. Ideally, such a resolution must also account for any apparent contradiction. If this is not done then individual investigators will construct radically different interpretations of the course of an incident. For instance, the NTSB report documents the Housing Authority employees' observations without attempting to resolve the apparent contradiction. As a result, it is impossible for readers to accurately assess whether or not a cursory visual inspection of the site should have identified the need for further support. In our example, some people will choose to follow the first account shown on the left of Figure 9.41. Others will choose to believe the alternative version of events shown on the right.

It is often impossible to entirely avoid ambiguity and inconsistency within an incident report. Many failures have complex organisational and managerial causes. These cannot easily be associated with discrete events that can be logged or recorded using automated equipment. Even when these devices are available, they often fail to provide unambiguous evidence. For example, many modern devices cannot record data at the same rate at which it is used by application processes [222]. Similarly, Chapter 5 has shown that the information provided by many of these recorders has been corrupted by reliability problems and design flaws. Even when accurate data is available, there can be genuine disagreement about the interpretation of that evidence. All of these factors make it unlikely that we shall have complete and unambiguous evidence for the events that contribute to major incidents. The key point, therefore, is not that the techniques in this chapter will entirely avoid ambiguity and inconsistency. They can, however, identify and address inconsistency *if investigators believe that it plays a significant role in our understanding of the incident.* In our example, the NTSB investigators did not further investigate the apparent contradiction between the Housing Authority employees and the other witnesses because even if the saw horses had remained in position they still failed to provide sufficient support for the exposed gas line.

Impact

The previous sections in this chapter have shown how a range of textual and graphical notations can be used to map out the events that contribute to safety-critical incidents. It has been argued that this form of modelling inevitably involves a process of selection or filtration. Secondary and primary investigations, typically, yield a mass of evidence about the course of an incident. Some of this evidence helps to establish the context in which a failure occurred. Other information provides more significant insights into the root causes of an incident. However, there will also be a mass of circumstantial data that has little apparent significance for the course of events. Investigators must, therefore, select which information is to be propagated into any reconstruction. For instance, the NTSB investigators gathered evidence about the excavation crews shift patterns immediately prior to the Allentown explosion. These were not found to have had any influence on this incident and so the information was not included in the time-lines and other reconstructions that were developed during the subsequent investigation.

The use of Petri Nets, of logic, or of Fault Trees only provides a crude indication of the salience of a particular event. The decision whether or not to include an event does not reflect the more detailed distinctions between root causes, contributory factors and contextual factors that were introduced in Chapter 7. This is a significant limitation. For instance, the NTSB summarised the outcome of their incident investigation in the following terms:

> "The National Transportation Safety Board determines that the probable cause of the natural gas explosion and fire at Gross Towers in Allentown, Pennsylvania, was the failure of the management of Environmental Preservation Associates, Inc., to ensure compliance with OSHA's and its own excavation requirements through project oversight. Contributing to the accident was the failure of the workmen from Environmental Preservation Associates, Inc., to notify UGI Utilities, Inc., that the line had been damaged and was unsupported.
>
> Contributing to the severity of the accident was the absence of an excess flow valve or a similar device, which could have rapidly stopped the flow of gas once the service line was ruptured. Also contributing to the severity of the accident was the absence of a gas detector, which could have alerted the fire department and residents promptly when escaping gas entered the building." [588]

Previous sections have not shown how such detailed assessments might be represented amongst the mass of events that we have represented in the previous Fault Trees, time-lines, Petri Nets and logic clauses. Before presenting one means of addressing this limitation, it is first important to clarify what we mean by terms such as 'root cause' or 'contributory factor'. The following list summarises the distinctions introduced in Chapter 7 but also introduces the term 'exacerbating factor'. This is identified in the NTSB conclusions and extends any impact analysis to consider events that occur in the immediate aftermath of an incident:

- *Contextual Factor.* Contextual factors are events or conditions that did not directly contribute to an incident.

- *Contributory Factor.* Contributory factors are events or conditions that collectively increase the likelihood of an accident but that individually would not lead to an adverse occurrence.

- *Root Cause.* Root causes capture Lewis' notion of causation established by counterfactual reasoning [491]. If a root cause had not occurred in the singular, particular causes of an incident then the incident would not have occurred.

- *Exacerbating factor.* Exacerbating factors do not contribute to the likelihood of an event but they can act to increase the consequences of an incident.

Figure 9.42 builds on Figure 9.11 to show how some of these distinctions might be represented within the fault tree notation. As can be seen, this embodies some of the NTSB investigators' findings, cited in the previous paragraph. This lack of training in OSHA excavation requirements is identified as a root cause for the incident. The fact that UGI were not informed that the line was uncovered is represented as a contributory factor.

Figure 9.42 again reflects the way in which simple syntactic extensions can be used to extend what can be represented in a modelling notation. However, it should be noted that we have only

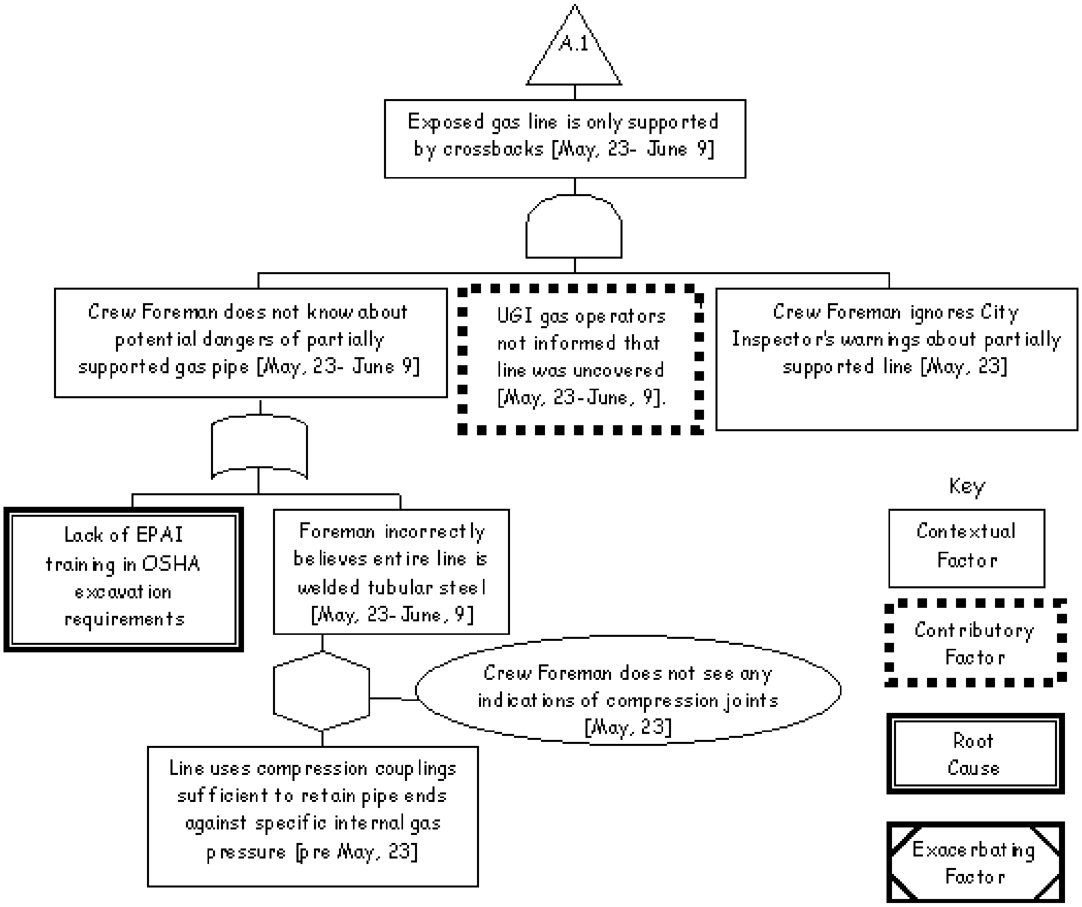
A.1
Exposed gas line is only supported by crossbacks [May, 23- June 9]
Crew Foreman does not know about potential dangers of partially supported gas pipe [May, 23- June 9]
UGI gas operators not informed that line was uncovered [May, 23-June, 9].
Crew Foreman ignores City Inspector's warnings about partially supported line [May, 23]
Lack of EPAI training in OSHA excavation requirements
Foreman incorrectly believes entire line is welded tubular steel [May, 23- June, 9]
Crew Foreman does not see any indications of compression joints [May, 23]
Line uses compression couplings sufficient to retain pipe ends against specific internal gas pressure [pre May, 23]
Key
Contextual Factor
Contributory Factor
Root Cause
Exacerbating Factor

Figure 9.42: Representing the Criticality of Distal Causes

provided an informal semantics for the different impact assessments that are represented in this picture. Similarly, we have not provided any grammatical rules that can be used to determine whether or not Figure 9.42 is well formed. For example, it might be argued that if a root cause is identified in a child event then that criticality should be propagated up the fault tree. By this argument, the intermediate event labelled Crew Foreman does not know about potential dangers of partially supported gas pipe should be denoted as a root cause that is inherited from the basic event labelled Lack of EPAI training in OSHA excavation requirements. We have chosen not to do this in order to keep Figure 9.42 as simple as possible. The ad hoc nature of these extensions re-iterates the point that we have used fault trees in a semi-formal manner. It would, of course, be possible to introduce mathematically defined rules to govern the representation of criticality within a fault tree. We have chosen not to do this. This decision is justified partly, as mentioned above, for the sake of simplicity. This decision is also justified by the relative lack of information that we have about the nature of criticality in incident investigations. We shall return to this theme in the next chapter. For now it is sufficient to observe that, in practice, it can be far harder to distinguish between root causes and contributory factors than might, at first, appear from Lewis' counterfactual definition.

Figure 9.42 is interesting for a number of reasons. Not only does it illustrate that impact or criticality assessments can be introduces as syntactic extensions to a semi-formal modelling notation, it also provides some insights into the Allentown incident. As can be seen, both the root cause and the contributory factor are identified as distal factors. In other words, they relate to events that occurred well before the gas leak or the explosion. In this respect, the NTSB investigators provide a good example of the 'systems' approach to incident investigation. They go beyond the immediate failures of individual staff to look at the longer term causes of the incident. This analysis can also be explained in terms of Mackie's ideas on particular and general causation. When attempting to assess criticality, there is a tendency for investigators to consider the general causes of an incident. In other words, the most significant or critical failures tend to be those that might threaten the safety of other applications rather than the particular failures associated with the incident under consideration.

The previous diagrams in this section have shown how impact assessments can be introduced into fault tree models. By denoting particular nodes as contributory factors or root causes, we have begun to indicate those events that might jeopardise the safety of future systems. It is important to emphasise that this involves a subjective classification. It reflects investigators' view of the relative criticality of key events during the course of an accident. However, it is important not to underestimate the importance of diagrams such as Figure 9.42. Too often these assessments are left as implicit judgements during the investigation process [427].

Figure 9.43 builds on the previous analysis by presenting an impact analysis of the proximal events that led to the Allentown incident. This diagram is based on the Fault Tree that was introduced in Figure 9.10. There are, however, two additions. The impact analysis was guided by the NTSB's findings, quoted above. As a result two additional events were incorporated into Figure 9.43. The first is labelled Excess flow valve not installed in Gross Towers. The second is labelled Gas detector capable of warning UGI was not installed in Gross Towers. These were identified as contributory factors by the NTSB but were not introduced into our model of the proximal events leading to the Allentown incident. This omission is very revealing. It emphasizes the way in which our initial model, represented by the Fault Tree in Figure 9.43, was initially constructed around those events that we knew to have taken place immediately before the explosion. The impact analysis, denoted by the fault tree in Figure 9.43, forced us to consider the way in which those events were affected by omissions or actions that did not take place. Several authors have noted that our bias in Figure 9.10 is symptomatic of a more general tendency to consider errors of commission rather than errors of omission [363].

Further biases affect the modelling and analysis of safety-critical incidents. We have already argued that there is a tendency to focus on contributory factors or root causes rather than the mitigating factors that help to reduce the consequences of an incident. This point can be illustrated by Figure 9.44, which is the same as Figure 9.13. The NTSB investigators focussed their analysis on the root causes and contributory factors that led to the incident. They did not devote the same amount of attention to the mitigating factors that contributed to the effective response after the Allentown explosion. In Mackie's terms, the investigation focuses on the general causes of failure. The investigators identified the particular events that occurred after this incident. In consequence, it can be difficult to identify the wider lessons that might be drawn from the successful response. This is worrying. Perrow and Sagan point to the difficulties of predicting future failures. We often fail to identify the general causes of particular incidents until a large number of similar failures have occurred. We might, therefore, learn more by studying an effective response than by trying to derive the general form of a particular failure.

This chapter has argued that primary and secondary investigations gather evidence about the events that are contribute to major failures. This evidence is then filtered to identify the key events that must be represented in any incident reconstruction. These models can then be used to distinguish root causes from other contributory and contextual factors. It is important to stress, however, that this only represents the first stage of analysis in any incident investigation. Previous paragraphs have re-iterated the problems that arise when attempting to derive general conclusions from the specific events that characterise a particular failure. It is, therefore, important that investigators can examine the products of such a generalisation to determine whether the wider conclusions accurately reflect their interpretation of the salient events that took place during an incident. Figure 9.45 represents one means of achieving this. The same fault tree notation is used to map out the conclusions of the NTSB report into the Allentown incident. As can be seen, this diagram avoids the timing information that was important in reconstructing the event-based models. Similarly, it omits some of the incidents that were considered to be significant in explaining the course of the Allentown incident but which are unlikely to recur in future failures. The detailed communications between the

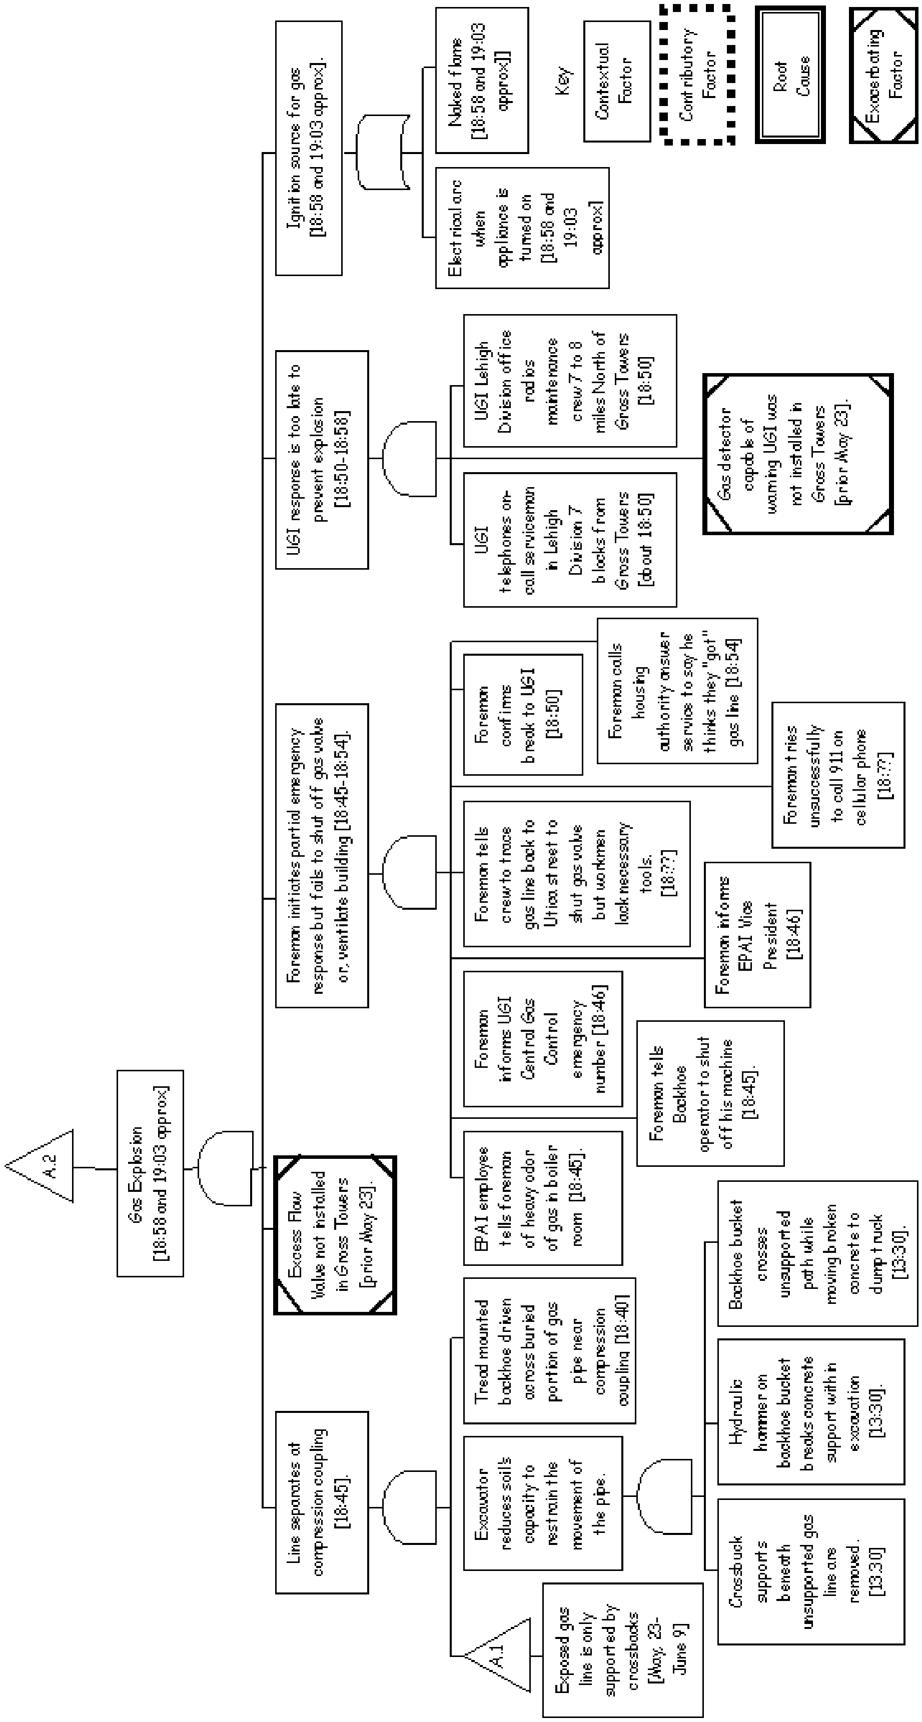
A.2
Gas Explosion [18:58 and 19:03 approx]
Excess Flow Valve not installed in Gross Towers [prior May 23].
Line separates at compression coupling [18:45].
Foreman initiates partial emergency response but fails to shut off gas valve or, ventilate building [18:45-18:54].
UGI response is too late to prevent explosion [18:50-18:58]
Ignition source for gas [18:58 and 19:03 approx].
Naked flame [18:58 and 19:03 approx]]
Electrical arc when appliance is turned on [18:58 and 19:03 approx]
Key
Contextual Factor
Contributory Factor
Root Cause
Exacerbating Factor
UGI Lehigh Division office radios maintenance crew 7 to 8 miles North of Gross Towers [18:50]
UGI telephones on-call serviceman in Lehigh Division 7 blocks from Gross Towers [about 18:50]
Gas detector capable of warning UGI was not installed in Gross Towers [prior May 23].
Foreman confirms break to UGI [18:50]
Foreman calls housing authority answer service to say he thinks they "got" gas line [18:54]
Foreman tries unsuccessfully to call 911 on cellular phone [18:??]
Foreman tells crew to trace gas line back to Utica street to shut gas valve but workmen lack necessary tools. [18:??]
Foreman informs EPAI Vice President [18:46]
Foreman informs UGI Central Gas Control emergency number [18:46]
Foreman tells Backhoe operator to shut off his machine [18:45].
EPAI employee tells foreman of heavy odor of gas in boiler room [18:45].
Tread mounted backhoe driven across buried portion of gas pipe near compression coupling [18:40]
Excavator reduces soils capacity to restrain the movement of the pipe.
Backhoe bucket crosses unsupported path while moving broken concrete to dump truck [13:30].
Hydraulic hammer on backhoe bucket breaks concrete support within excavation [13:30].
Crossbuck supports beneath unsupported gas line are removed. [13:30]
A.1
Exposed gas line is only supported by crossbucks [May, 23-June 9]

Figure 9.43: Representing the Impact of Proximal Causes

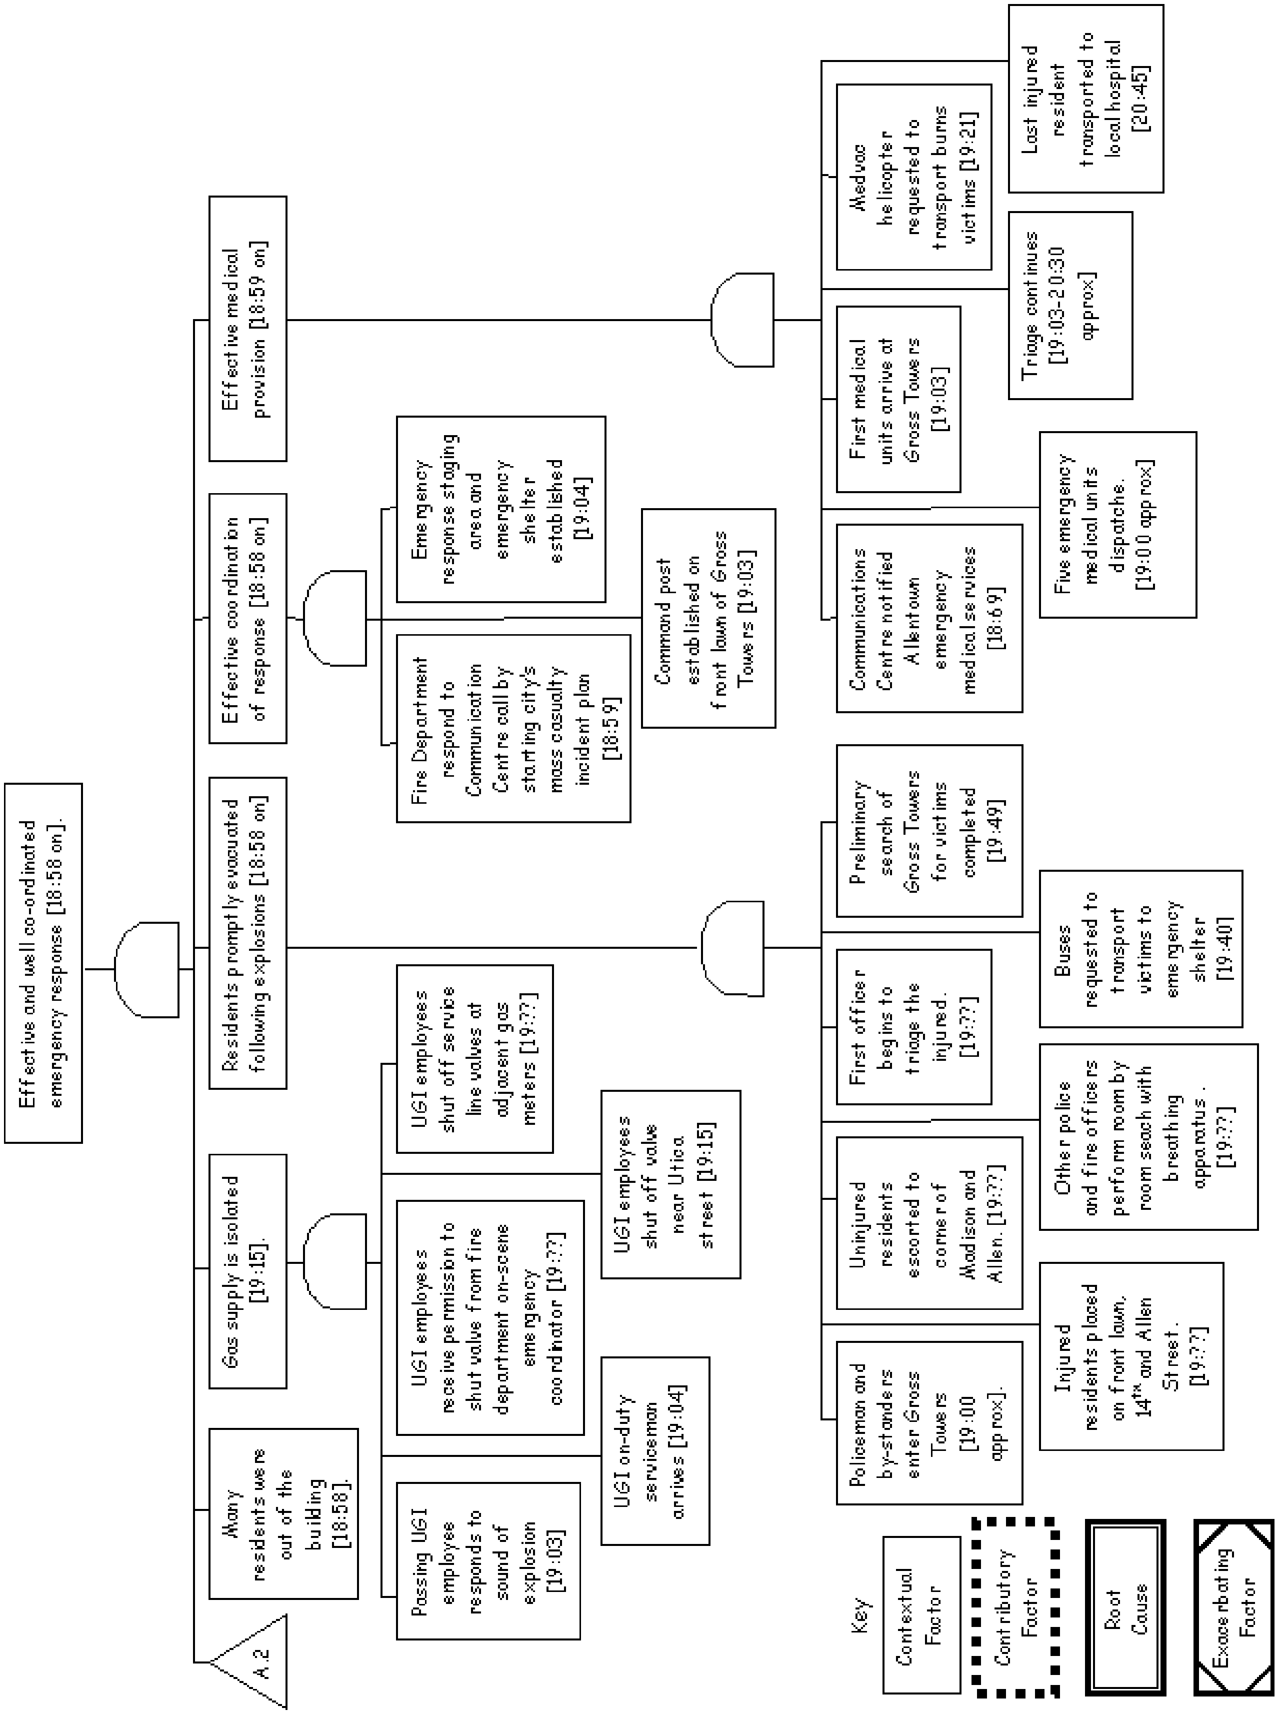
Effective and well co-ordinated emergency response [18:58 on].
A.2
Many residents were out of the building [18:58].
Gas supply is isolated [19:15].
Residents promptly evacuated following explosions [18:58 on]
Effective coordination of response [18:58 on]
Effective medical provision [18:59 on]
Passing UGI employee responds to sound of explosion [19:03]
UGI on-duty serviceman arrives [19:04]
UGI employees receive permission to shut valve from fire department on-scene emergency coordinator [19:??]
UGI employees shut off valve near Utica street [19:15]
UGI employees shut off service line valves at adjacent gas meters [19:??]
Fire Department respond to Communication Centre call by starting city's mass casualty incident plan [18:59]
Command post established on front lawn of Gross Towers [19:03]
Emergency response staging area and emergency shelter established [19:04]
Policeman and by-standers enter Gross Towers [19:00 approx].
Injured residents placed on front lawn, 14th and Allen Street. [19:??]
Uninjured residents escorted to corner of Madison and Allen. [19:??]
Other police and fire officers perform room by room seach with breathing apparatus. [19:??]
First officer begins to triage the injured. [19:??]
Buses requested to transport victims to emergency shelter [19:40]
Preliminary search of Gross Towers for victims completed [19:49]
Communications Centre notified Allentown emergency medical services [18:69]
Five emergency medical units dispatche. [19:00 approx]
First medical units arrive at Gross Towers [19:03]
Triage continues [19:03-20:30 approx]
Medvac helicopter requested to transport burns victims [19:21]
Last injured resident transported to local hospital [20:45]
Key
Contextual Factor
Contributory Factor
Root Cause
Exacerbating Factor

Figure 9.44: Representing the Impact of Mitigating Factors

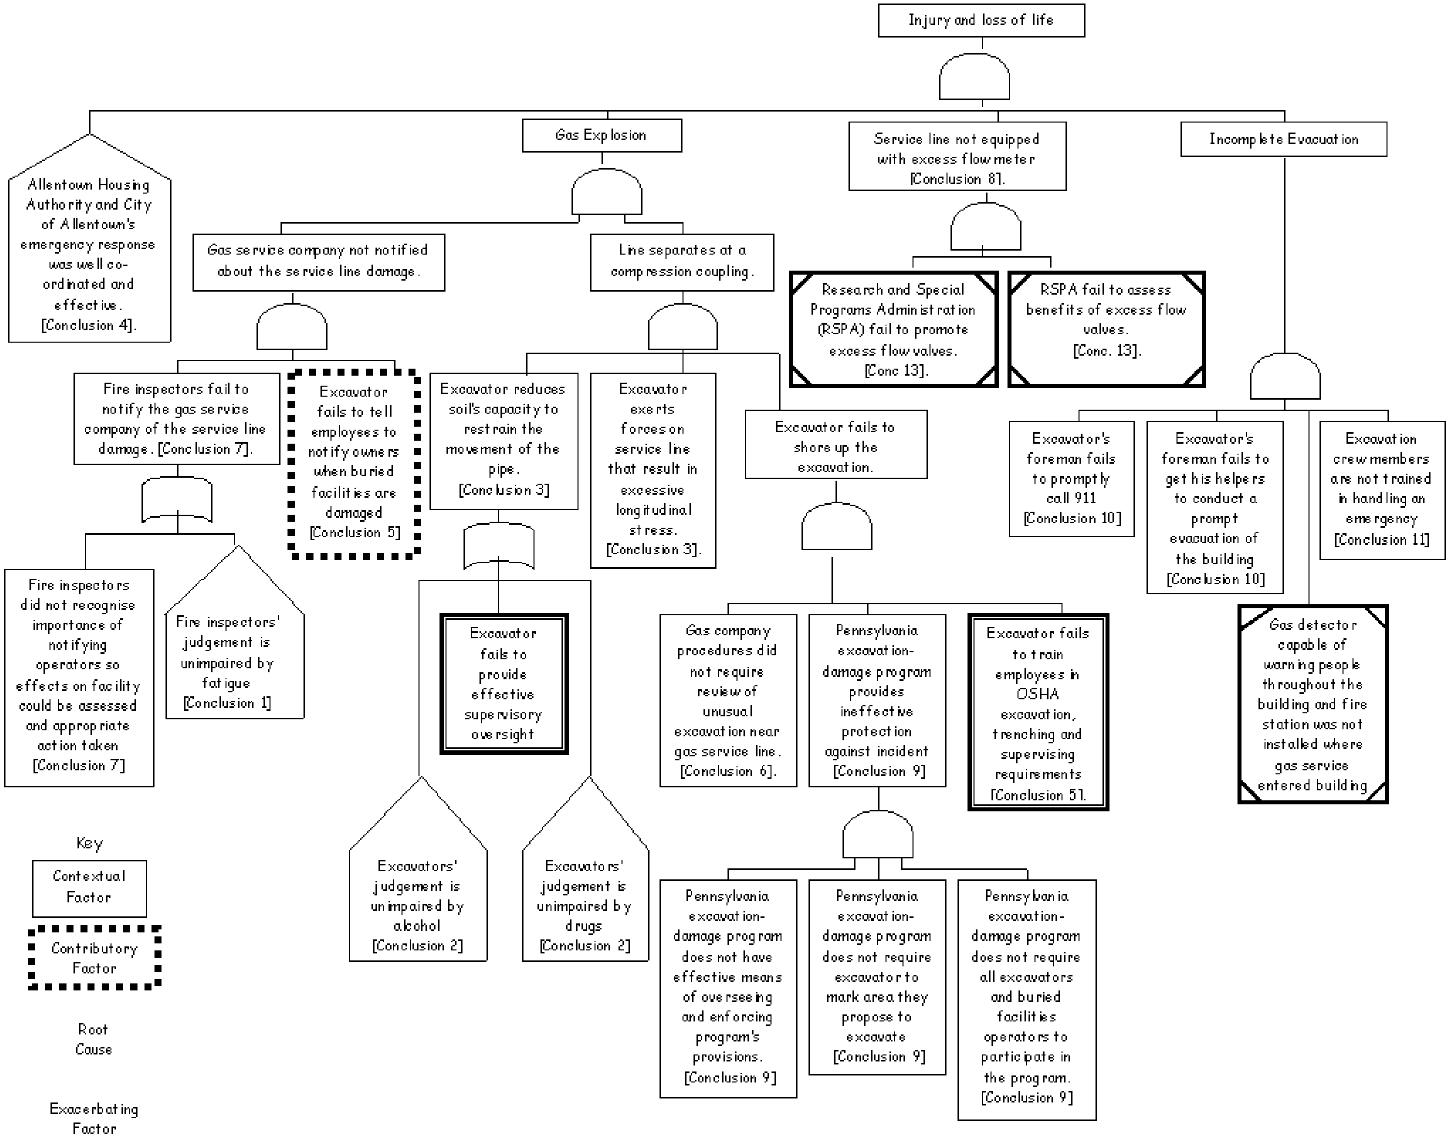
Injury and loss of life
Gas Explosion
Service line not equipped with excess flow meter [Conclusion 8].
Incomplete Evacuation
Allentown Housing Authority and City of Allentown's emergency response was well co-ordinated and effective. [Conclusion 4].
Gas service company not notified about the service line damage.
Line separates at a compression coupling.
Research and Special Programs Administration (RSPA) fail to promote excess flow valves. [Conc 13].
RSPA fail to assess benefits of excess flow valves. [Conc. 13].
Fire inspectors fail to notify the gas service company of the service line damage. [Conclusion 7].
Excavator fails to tell employees to notify owners when buried facilities are damaged [Conclusion 5]
Excavator reduces soil's capacity to restrain the movement of the pipe. [Conclusion 3]
Excavator exerts forces on service line that result in excessive longitudinal stress. [Conclusion 3].
Excavator fails to shore up the excavation.
Excavator's foreman fails to promptly call 911 [Conclusion 10]
Excavator's foreman fails to get his helpers to conduct a prompt evacuation of the building [Conclusion 10]
Excavation crew members are not trained in handling an emergency [Conclusion 11]
Fire inspectors did not recognise importance of notifying operators so effects on facility could be assessed and appropriate action taken [Conclusion 7]
Fire inspectors' judgement is unimpaired by fatigue [Conclusion 1]
Excavator fails to provide effective supervisory oversight
Gas company procedures did not require review of unusual excavation near gas service line. [Conclusion 6].
Pennsylvania excavation-damage program provides ineffective protection against incident [Conclusion 9]
Excavator fails to train employees in OSHA excavation, trenching and supervising requirements [Conclusion 5].
Gas detector capable of warning people throughout the building and fire station was not installed where gas service entered building
Excavators' judgement is unimpaired by alcohol [Conclusion 2]
Excavators' judgement is unimpaired by drugs [Conclusion 2]
Pennsylvania excavation-damage program does not have effective means of overseeing and enforcing program's provisions. [Conclusion 9]
Pennsylvania excavation-damage program does not require excavator to mark area they propose to excavate [Conclusion 9]
Pennsylvania excavation-damage program does not require all excavators and buried facilities operators to participate in the program. [Conclusion 9]
Key
Contextual Factor
Contributory Factor
Root Cause
Exacerbating Factor

Figure 9.45: Representing Impact in a Causal Analysis

Fire Service coordinator and the UGI employees is an example of such a particular event. Although this Fault Tree captures the more general conclusions about this incident, it is still possible to distinguish those findings that relate to the root cause from those that relate to contributory factors and findings that relate to the contextual factors from those that relate to exacerbating/mitigating factors. This illustrates how the same graphical notation can be adapted to support the transition between incident modelling, which was the focus of this chapter, and causal analysis, which is the focus of the next chapter.

9.3 Summary

This chapter has introduced a number of modelling notations, which can be used to reconstruct the events that lead to safety-related incidents. These languages help to strip out the clutter of contextual information that threatens to obscure important information about adverse occurrences. They can chart proximal and distal failures so that investigators can establish both the immediate and longer term events that contribute to an incident. They provide an overview of the interaction between human, technical and organisational failures. This is important because this diverse range of events cannot easily be represented within many of the simulation environments introduced in Chapter 8. Reasoning and proof techniques can also be used to check for the consistency and completeness of the resulting models. Reconstruction techniques, therefore, help to develop coherent accounts from the diverse evidence that is elicited during primary and secondary investigations. These reconstructions of the events leading to an incident can, in turn, be used to support hypotheses about the causes of incidents and accidents.

We have focussed on reconstructing events that contribute to an incident. It is important, however, to represent both the commission of undesirable events as well as the omission of necessary actions. The dual nature of any reconstruction has not been stressed enough in the preceding discussion. This is partly due to the nature of the Allentown incident. The NTSB team focussed on those actions that actively contributed to the explosion. OSHA conducted a separate investigation into those procedures and guidelines that were ignored during the EPAI excavations. This ultimately led to a Citation and Notification Penalty for approximately $54,000. If we had focussed on reconstructing the incident from OSHA's viewpoint then these omissions would have formed a far more significant component of the model. This illustrates another important point. Reconstructions focus on critical events during an adverse occurrence. The exact definition of what does and what does not constitute a 'critical' event is determined by the person building the model. The focus of the NTSB investigation was clearly different from that conducted by OSHA's employees and hence we would expect some important differences between the reconstructions that they might develop. However, if we could develop some common tools and techniques these is the possibility that future investigations might share reconstructions to support these different forms of analysis.

The development of an incident reconstruction is not an end in itself. The utility of any notation is determined by whether or not groups of individuals can use that notation to cooperate on the development of a natural language, accident report. This raises a number of further issues. The first set of problems relate to the difficulty of constructing coherent temporal models for safety-related incidents. It is a non-trivial task to resolve the contradictory timings that often appear in eye-witness evidence and automated logs. It can also be difficult to integrate imprecise temporal information about operator behaviour with the more precise temporal schemas that are available for process components. It is important to stress that the development of coherent temporal models must not force analysts into arbitrary decisions or commitments to timings that are not supported by the available evidence.

It is not simply important that a reconstruction notation is capable of representing the course of events, it is also important that investigators can learn to exploit those capabilities. We have argued that there is often a trade-off between the visual appeal of formal and semi-formal notations and the reasoning power that those notations offer to analysts and investigators. This is significant because formal proof techniques provide a powerful means of identifying the temporal ambiguities that have been criticised in the previous paragraph. Tool support has been identified as one means of improving the ‘usability’ of notations with a relatively low visual appeal. However, further work is urgently required to determine whether similar tools, that have been developed in other areas of engineering, can be applied to analyse incident reconstructions.

The final set of problems stem from the difficulties of managing cooperative work between heterogenous groups of experts. Rather than focusing on modelling capabilities or visual appeal, these problems relate to aspects of control. For example, what are the consequences of allowing more than one author to simultaneously work on a formal or semi-formal description of an incident? No research has been done into these issues. This is an important omission. Without some understanding of the group processes involved in incident reconstruction, it is unlikely that adequate tool support can be developed. This may explain why many existing systems, such as Fault-Tree editors, often only support specific areas of an investigation. They are frequently restricted to systems or control flow analysis. Few attempts have been made to support human factors investigations or the analysis of managerial decision making.

We have focussed on graphical and textual time-lines, on fault trees and Petri Nets and on temporal extensions to first order logic. It is important to emphasise that these represent a very small subset of the range of notations that are currently being applied to this area. For example, we have cited work into more complex logics that include explicit notions of causation [470] or obligation and permission [118]. Others have used state-based techniques that are amenable to model checking [192]. It is too early to judge which, if any, of these approaches will be accepted by practitioners. However, the increasing complexity of many technological failures makes it highly likely that the incident investigators of the future will have to exploit more formal techniques for incident reconstruction.

Chapter 10

Causal Analysis

This book is based around an implicit model of incident reporting. The evidence that is collected during primary and secondary investigation helps to reconstruct the event leading to an adverse occurrence. The resulting models and simulations can then analysed be to distinguish root causes from contributory factors and contextual details. Previous chapters have briefly introduced the analytical techniques that can be used to identify the most salient events from a more general reconstruction. The following pages build on this by describing the aims and objectives of such techniques in more detail.

10.1 Introduction

Chapters 8 and 9 have described how simulation and modelling techniques can be used to reconstruct the events that lead to failure. Causal analysis looks beyond *what* happened to identify the reasons *why* [248]. Kjellén [444] identifies three broad approaches to causal analysis:

- *Expert judgement.* Even with the support of analytical and statistical techniques, mentioned below, it is difficult to prevent investigators from forming subjective judgements that help to shape and direct the causal analysis of any incident. These judgements influence every stage of the investigatory process and so can have a profound impact upon the nature and extent of the evidence that is obtained before any causal analysis even begins. It is important to emphasise that subjective judgements need not be 'a bad thing'. They reflect the expertise and experience of the investigator. As we shall see, many of the recommended analytical and statistical techniques do little more than document the process of forming these judegements so that they are open to challenge, or validation, through peer review;

- *Statistical techniques.* These techniques are, typically, applied to identify common causal factors amongst a collection of similar incidents. They help to determine whether the presence or absence of certain factors increases the probability that an incident will occur. At its simplest, statistical techniques can track uni and bi-variate distributions [471]. Chapter 3 has, however, argued that many incidents have complex, inter-connected causes. Chapter 5 has also argued that the limitations of automated logging systems and the unreliability of human witnesses can filter the evidence that is obtained about these causal factors. Some researchers have, therefore, begun to explore multi-variate techniques. There have also been some initial attempts to exploit Bayes' theorem as a means of quantifying the likelihood that a particular root cause led to an incident given particular observations of that incident. This work builds upon attempts to assess the reliability of software systems given uncertain information about potential failure modes [497]. These techniques can also be used post hoc to parameterise expert assessments about the likely causes of an incident. Chapter 15 will describe this work in more detail. It will also describe the practical limitations that have restrict the application of these more advanced statistical techniques;

- *Analytical techniques.* These techniques provide a broad range of formal and semi-formal techniques that are inteneded to support causal analysis. Many of these approaches rely upon counterfactual reasoning. Causal factors can be distinguished from contributory factors and contextual details if it can be argued that if the causal factor had not occurred then the incident would not have occurred. As we shall see, it can be difficult to apply this form of reasoning to certain incidents especially when the failure of a barrier is identified as a potential causal factor. Other analytical techniques, therefore, rely upon checklists. Investigators are guided in their causal analysis by a limited number of pre-defined categories that help to identify common factors in previous incidents. These approaches are limited in that investigators may be biased towards particular categores, for example those that appear at the top of a list. This can hinder a more coherent causal analysis.

Chapter 11 surveys a number of different statistical and analytical techniques. Chapter 15 also provides an overview of the use of statistical techniques in monitoring the changing causes of incidents between industries and within the different groups of similar organisations. In contrast, this chapter focusses on the detailed application of one particular set of analytic tools. The strength and weaknesses of techniques advoctaed by NASA [571] and by the US Department of Energy [209, 208] are demonstrated using the loss of the Mars Climate Orbiter and Mars Polar Lander. Neither of these case studies is 'safety-critical'. They have, however, been chosen because they illustrate the general applicability of incident reporting techniques to investigate the failure of dependable systems. These two concepts are closely related [850]. Their similarities can be used to advantage of borrowing techniques from one to deal with the other [486]. The NASA case studies were also chosen because of the technological sophistication of the systems involved, they therefore represent a strong contrast with the National Transportation Safety Board (NTSB)'s Allentown incident in Chapter 9.

It can, in practice, be difficult to distinguish between the stages of investigation, reconstruction and analysis. Investigators may be forced to obtain more evidence to resolve the omissions and ambiguities that are identified when they reconstruct the events leading to failure [462]. Similarly, investigators often have to extend the scope a reconstruction as new theories are developed about the cause of an incident. Chapter 6 has also described how the collection of evidence can be biased, or 'focussed', by an investigator's working hypotheses about the probable course of events. These pragmatic issues can complicate the application of the modelling techniques that have been introduced in previous chapters. The costs associated with the development of interactive three-dimensional simulations can dissuade investigators from revising them in the light of new causal hypotheses. Similarly, the problems of maintaining large and complex graphical models can force investigators to use techniques that have stable tool support. The closing sections of this chapter, therefore, attempt to assess the practical implications of the analytical techniques that are introduced. In particular, there is a concern to assess the degree to which these approaches support 'real world' investigation practices [73].

10.1.1 Why Bother With Causal Analysis?

Incident analysis techniques, typically, provide means of distinguishing root causes from contributory factors and contextual details. Chapter 7 introduced these different causal concepts. They can be summarised as follows. A causal factor was described using a counterfactual argument [491]. If a causal factor had not occurred then the incident would not have occurred. If A and B are states or events, then A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either. It is important to emphasise that this is based on Mackie's idea of singular causality [508]. Singular causality is used because there may be other failures that could have had the same consequences but which did not occur in this instance. In contrast, root causes depend upon a more general view of causality. These are causes that have the potential to threaten the safety of future systems. They may, in turn, contribute to a number of the causal factors that are observed in a particular incident. In contrast, contributory factors can be thought of as individually necessary but not globally sufficient [677]. These are events or conditions that collectively increase the likelihood of an accident but that would not themselves lead to an adverse occurrence. Finally, contextual details are events or conditions that did not directly contribute to an incident. They help to set the scene and establish the context in which an adverse occurrence took place. They can also help to establish that certain factors were NOT significant in the events leading to failure.

It might seem superfluous to ask why analytical techniques have been developed to distinguish between the factors described in the previous paragraph. It is clearly important to analyse the circumstances of a near miss to determine how best to avoid any recurrence that might result in more severe consequences. Within this high level goal, there are a number of more detail motivations for incident analysis. These different motivational factors can have an important effect in determining which analytical techniques will offer the greatest benefits for any particular organisations. These justifications for incident analysis can be summarised as follows:

- *analysis is a regulatory requirement.* In many industries, organisations must analyse their incident reports in order to meet regulatory requirements. For example, ICAO Annex 13 requires that member states not only analyse the causes of individual aviation incidents but also that organisations must use this analysis to identify any common causes between similar reports [384]. Similarly, the UK Rail Inspectorate's assessment criteria for safety cases requires that all operators demonstrate "established adequate arrangements for identifying the causes of incidents" [350]. Even if there is no regulatory requirement, institutional and organisational policy often requires that a causal analysis should be performed. For instance, the US Army has published detailed recommendations that can be used to determine potential causal factors during an incident investigation [803]. NASA have published similar guidelines [571].

- *analysis is a prerequisite for statistical comparisons.* Regulators are concerned to ensure that organisations identify the causes of potential incidents. This is important if companies are to learn from previous failures. Companies must also analyse the causes of potential incidents because regulators use this information to target their intervention in the market place. Causal information from individual companies is, typically, entered into a central database. This database is then queries at regular intervals to identify common causal factors and also to generate a 'most wanted' list of safety improvements within an industry. The UK Health and Safety Executive recently announced its initiative reduce the fatality and major injury rate from 260 per 100,000 workers in 1999/2000 to 230 per 100,000 workers by 2009/2010. Together with these targets they have also announced a review of their incident reporting regulations [335]. The HSE recognise that the overall effectiveness of any safety intervention is determined by the regulator's ability to identify the root causes of common incidents. The review indicates the need to have confidence in the analytical and reporting procedures that inform each statistical return.

- *focus for remedial actions.* The most immediate reason for performing a causal analysis is to focus remedial actions in the aftermath of an incident. Short-term resources should address the root cause before any contributory factors. Once investigators have addressed immediate concerns over the root cause of an incident, additional resources can be allocated to other events and conditions that contributed to the incident. It is apparent, however, that any disagreement about the causes of an adverse occurrence can have profound consequences. Similarly, significant problems can arise if the analysis fails to correctly identify the root cause of an incident. Under such circumstances, the investigators' ability to prevent a potential recurrence will be compromised by the allocation of resources to less significant aspects of a system. This is illustrated by the way in which poor training is often identified as a root cause of medical incidents rather than the poorly designed equipment and long working hours that staff are forced to endure [121].

- *guiding the allocation of development resources.* At an organisational level, incident reporting schemes are often argued to be an effective means of informing risk analysis. As we shall see, however, many organisations do root cause analysis but do not feed the data into design. Information about previous failures can be used to direct both acquisition and development work. Such an integrated approach can only be successful if organisations can correctly identify those components and processes that contributed most to an incident. If the analysis of an adverse occurrence is biased by political or organisational pressures then there is a danger that other aspects of a system will be unnecessarily implicated in the causes of an incident. Long-term development resources may allocated to areas that do not pose the greatest threat to future incidents. This is illustrated by the Fennell report which argues that the London Underground Management "...remained of the view that fires were inevitable in the oldest most extensive underground system in the world" [247]. The root cause of these fires, in particular the built up of detritus in key areas of the system, was not addressed. Instead, staff were trained to detect and respond to these incidents once they had started. There continued to be a steady number of minor fires until the Kings Cross' accident.

- *characterisation of causal complexes.* The causal analysis of incidents need not simply focus on identifying a single root cause. This has been a weakness in the statistical returns that have been required by some regulators. As many authors have observed, incidents and accidents typically stem from pathological combinations of events [699]. As much can be learned from the ways in which those failures combine as can be learned from single causal factors in isolation. This poses a number of problems. Rather than describing safety priorities in terms of a 'hit list' of individual causal factors, it may be more important to identify critical patterns of events. For example, the recruitment of a new sub-contractor followed by a component failure or the installation of a new item of equipment shortly before a software release. It is for this reason that many organisations, including the European Space Agency and the US Navy [5], have begun to look beyond simple categorisations of causal factors. Later sections will describe this 'lessons learned' work in more detail. For now, however, it is sufficient to observe that they have developed data mining and information retrieval techniques that help investigators to identify patterns within a collection of previous incidents [413].

These motivations provide criteria that can be used to assess the utility of different analysis techniques. For example, the previous chapter briefly explained how the minimal cut set of a fault tree can be used to support incident analysis. The elements of this set represent the smallest possible conjunction of events in which if any basic event is removed then the top condition will not occur [27]. Root causes are basic events that are common to every member of the minimal cut set. There is no reason why there should not be multiple root causes that are common to the elements of this set. In consequence, this approach cannot easily be applied to identify a unique root causes.

There are further tensions between the different motivations that support the causal analysis of near miss incidents. As we shall see, some analytical techniques identify a 'primary causal factor'. These techniques, typically, require that investigators select the most significant cause from a predetermine list of potential factors. This approach helps to ensure consistency between different investigators. The use of an agreed list helps investigators to avoid using a range of different terms to describe the same causal factors. This can, in turn, increase confidence in regulatory statistics. There are, however, a range of problems. It can be difficult to construct an appropriate list of agreed causal factors. As we have seen, new causal factors can emerge with the introduction of novel equipment and working practices. It can also be difficult to identify a single 'main' cause from many competing alternatives. Previous sections have shown how a single event can have multiple proximal and distal causes. Any one of these could be regarded as a root cause on the basis of Lewis' counterfactual arguments. For example, the Allentown incident might have been avoided if excess flow valves had been installed or if proper excavation procedures had been followed. Which of these is the true 'primary' cause?

This analysis illustrates a number of points that will be reiterated throughout this chapter. Firstly, analytical techniques must often be refined to support particular organisational objectives. For example, investigators are often expected to translate their findings into a form that is acceptable to regulatory organisations. This can involve the selection of a primary causal factor from an 'accepted' list of root causes. There is a danger that such requirements may prevent investigators from adequately considering the complex causes of many technological failures [675]. Secondly, causal analysis can yield important information for the subsequent development of safety-critical applications. It is, therefore, important that the products of such an analysis should be in a form that is compatible with subsequent risk assessment procedures. This does not imply that similar techniques should be used for both activities. However, it is important that designers can understand the outcome of any causal analysis. Finally, the term ‘causal analysis’ applies at several different levels. The previous discussion has used it to describe the process by which the root causes of a particular incident can be distinguished from contributory factors and contextual details. However, causal analysis can also be applied over collections of incidents. This is essential if investigators are to identify patterns of failure and emerging trends in a number of similar incidents.

10.1.2 Potential Pitfalls

Previous paragraphs have introduce some of the complexities that affect the causal analysis of adverse incidents. For example, regulatory requirements impose additional constraints upon the causal analysis of some incidents. The format that best supports ‘organisational learning’ may not be the best format to support the statistical analyses demanded by regulators. There are further complexities. In particular, analysts may lack the evidence that is necessary to perform a detailed causal analysis. Later sections will describe how design decisions and budgetary constrained determine that NASA’s Mars Polar lander would not provide any telemetry data during the Entry, Descent and Landing phase of the mission. In consequence, it was impossible for investigators to accurately reconstruct the events that led to the failure nor could they identify definitive root causes. The following paragraphs, therefore, examine further problems that can complicate the analysis of ‘near miss’ incidents:

- *The scope of a reporting system influences the scope of any causal analysis.* In an ideal situation, investigators would conduct an analysis in an environment that is free from external or organisational constraints. Unfortunately, this does not reflect the experience of most operational reporting systems. For example, local schemes deliberately restrict the scope of the investigator’s analysis to ‘target the doable’. Many hospital reporting systems identify failures within a particular department or ward [119]. They explicitly exclude reports that deal with failures in other departments or at higher levels in the management structure. This pragmatism effectively restricts the scope of any analysis to the immediate activities of the group that participates in the reporting scheme. Of course, the scope of any analysis can be widened as reporting systems are extended nationally and across an entire industry. In consequence, national and international reporting systems are being developed within the healthcare industry. However, these initiatives also place either explicit or implicit boundaries on the scope of any investigation. For example, the ASRS was deliberately established to cut across the many different professional and organisational demarkations that characterise the US aviation industry. It solicits input from commercial, military and general pilots. It encourages reports from air traffic controllers and ground staff. It is important to remember, however, that even this scheme is bounded by organisational factors. For instance, the ASRS provides relatively few insights into ‘near miss’ incidents involving military aircraft. This partly stems from a noticeable under-reporting, mentioned in Chapter 5. It also arguably reflects the ASRS’ analytical focus on commercial and general aviation.

- *Organisational factors place unnecessary constraints upon causal analysis.* Organisational goals and priorities influence any causal analysis. These influences do not simply act upon the individuals who report adverse occurrences. They must also affect incident investigators. The most obvious manifestation of this is the lack of critical analysis about regulatory intervention. As noted in the opening chapters, regulators are ultimately responsible for the safety record in most industries. Very few investigators ever analyse the impact that these organisations have upon the course of an incident. There are some notable exceptions to this, including the NTSB’s Allentown report that was cited in the previous chapter [588]. These exceptions, typically, occur when investigators are independent both from the regulator and from any organisation that is directly implicated in an incident. In particular, regulatory failure is most often exposed at the large scale public enquiries that follow major accidents [193]. Given the pragmatics of most reporting systems, it should not be surprising that such causal factors are not more apparent in the analysis of ‘near miss’ incidents.

- *Organisational can inform a causal analysis.* The previous paragraphs have stressed the way in which organisational factors can constrain the scope of any causal analysis. It is also important to emphasise that these factors can play a positive role. In particular, the last decade has seen a movement away from individual blame as a satisfactory causal interpretation of many incidents. This movement has been promoted by many researchers [701, 844]. However, their work would have had little weight if commercial and regulatory organisations had not had the insight to act upon it. In particular, it is important not to underestimate the powerful normalising influence that investigator training can have upon the products of any causal analysis. This can be seen in the impact of Crew Resource Management training in the aviation industry. This has equipped investigators with a vocabulary that can be used to describe the causes of failure in team-based communication and decision making. Before the widespread introduction of this training, investigators failed to derive many insights about the role of team factors in the causes of many incidents and accidents [57, 734, 410].

- *Historical factors help to shape any causal analysis.* The previous paragraph has argued that explicit training can inform an investigators' interpretation of the events leading to an incident. Implicit forms of training also play an important role in determining the outcome of any causal analysis. For instance, traditions of interpretation can become established within groups or teams of investigators. This can be seen as a strength; similar incidents are handled in a consistent manner. There is, however, a danger that investigators will become habituated to some causal factors so that they are identified irrespective of the circumstances surrounding a particular incident. In the past, human error was often seen as a routine cause of many incidents [718]. Increasingly, however, software is being identified as the predominant cause of many safety-critical incidents and accidents [411]. For example, later sections will describe the software failures that led to the loss of NASA's Mars Climate Orbiter and to difficulties in the Stardust programme. These failures clearly helped to focus the investigators attention on software failure as a potential factor in the subsequent loss of the Mars Polar Lander. It is important that the causes of previous incidents inform rather than bias subsequent investigations. This narrow distinction raises important pragmatic problems for investigators who must retain an open mind when they deploy finite analytical resources.

- *Causal analysis is constrained by available resources.* The second half of this chapter will present a range of analytical techniques that investigators can use to distinguish root causes from contributory factors and contextual details. These approaches differ in terms of the amount of time that investigators must invest before they can learn how to exploit them. They also offer different levels of tool support. These factors can have a profound impact upon which analytical techniques are chosen within a particular organisation. More complex techniques are less likely to be used in local reporting system that must rely upon the enthusiasm of key individuals with limited training in incident analysis. Resource constraint also affect national and regional systems. Investigators must justify resource expenditure to upper levels of management if they are to ensure continued support for a reporting system. This topic is addressed in the final chapters of this book. As we shall see, it is difficult to underestimate the importance of these cost-benefit decisions. Complex techniques will fail to provide analytical insights if they are under-resources. Conversely, these more advanced approaches often carry a significant overhead in terms of staff time that cannot be justified for many relatively simple incidents. However, it is equally important to emphasise that 'low-cost' analytical techniques often yield superficial results when they are applied to more complex incidents. The problem of selecting an appropriate analytical technique is compounded by the lack of empirical evidence, or published practical experience, that compares the costs and benefits of different forms of causal analysis.

- *Who Performs the Analysis?* The previous paragraphs provide an insight into the complexities that surround any causal analysis of adverse occurrences. As can be seen, many of these issues focus upon the organisational biases that affect any investigation. These biases can have both positive and negative influences with respect to the overall safety of an application. For instance, an emphasis away from individual error can be beneficial in encouraging investigators to look for wider causes of adverse occurrences. Similarly, by focusing on the 'doable' investigators can maximise the allocation of their finite resources. Organisational factors have a negative impact if individual or group objectives are considered to be more important than the overall safety of an application. It is for this reason that many reporting schemes rely upon outside organisations to analyse the reports that they receive. For example, the University of Strathclyde coordinates the analysis of incident data on UK Railways [197]. The ASRS is operated by Batelle under contrast from NASA. These external organisations assume responsibility for the analytical techniques that are then applied to each report. This approach has the benefit that investigators are seen to be independent from the organisations who must act on any recommendations. In practice, however, there remain strong implicit constraints on the forms of analysis that are performed even by external investigators. For example, a semi-competitive tendering process is often used to award the contracts for these systems. This process can focus the attention of the existing contract holder. It can also introduce terms of reference within a contract that place specific bounds on the form of analysis that is to be performed.

- *The Importance of Balancing Domain Expertise and Multi-Modal Skills.* The emergence of national and international systems has seen a new generation of professional incident investigators. These analysts fall into one of two categories. Firstly, domain specialists often 'move' into incident investigation after lengthy periods of field service. There are strengths and weaknesses to both approaches. Domain specialists can quickly lose touch with current operating practices in rapidly changing industries. In consequence, they must either undergo continual retraining to reinforce their existing skills or they must gather new ones. In particular, domain specialists often lack expertise in the human factors domain, they may also have little first hand experience of systems engineering. This makes their analysis vulnerable to criticisms from individuals with these more specialist skills. Secondly, there is a growing number of incident investigators who are recruited in spite of their lack of domain skills. These individuals contribute what can be termed 'multi-modal' analytical techniques. They provide tools from other engineering disciplines, such as human factors and systems engineering, that can be applied to analyse incidents in many different application domains. The situation is then reversed, the analytical insights provided by these individuals is then vulnerable to criticism by those who have first hand experience of the application domain. Such observations should emphasise the political nature of many investigations; there is a danger that any analysis may be jeopardised by disagreements between domain specialists and expert witnesses who possess these multi-modal skills. Some organisations, notably the Australian Transportation Safety Bureau , have launched a series of initiatives that are intended to find some middle ground [49]. They have deliberately distinguished between multi-modal and industry specific training requirements. Investigators from each mode of transportation are expected to possess these multi-modal skills, including human factors and systems engineering expertise. In addition, they must refresh the technical and practical foundations of their domain knowledge. However, the ATSB intend that their inspectors will be qualified in more than one domain. This will help to transfer multi-modal analytical techniques between road, rail, maritime and aviation investigations. Just as the US NTSB have established a reputation for their innovative use of simulation and reconstruction techniques, the ATSB continue innovate in the way that they train and deploy their investigators. It remains to be seen whether this transition from a narrow focus on domain expertise to a multi-modal approach will have a lasting impact on the nature of incident analysis within each mode of transportation.

- *The Importance of Justifying Causal Analysis.* The mutual vulnerability of domain specialists and multi-modal investigators raises a number of important concerns about the application of analytical techniques within many investigations. In particular, the individual investigator's interpretation of an incident is open to many different challenges. It is, therefore, very important that sufficient evidence is provided about the analytical techniques that are used to support the findings of any investigation. This has been a particular weakness of investigations into human factors issues. Frequently investigators refer to problems of high workload and poor situation wareness without explaining the particular observations that support these conclusions [408]. Of course, as noted above, not all of these analyses were performed by investigators with the relevant human factors training. Similar weaknesses can also be found in systems engineering accounts. For example, it is often difficult to replicate the vibrations that metallurgists have identified as a primary cause of metal fatigue in aircraft components. The ambivalent results of airborne and ground tests are occasionally omitted. In other instances, investigators place sparse details of negative results in appendices that are not then distributed with the body of a report. It can be argued that these techniques support the dissemination of important safety information. Most readers are unconcerned with the methods that were used to reach a particular conclusion. However, these same techniques can be viewed as rhetorical devices. The lack of analytical detail prevents other investigators from raising detailed objections to an analysts findings. It is for this reason that I believe all investigators should provide detailed documentation to support the findings of any analytical technique.

- *Avoid the over-interpretation of sparse data.* There are many reasons why investigators must document and justify their use of analytical techniques. In particular, there is a danger that individuals will be tempted to form conclusions that are not warranted by the evidence that is available. This tendency can be exacerbated by some of the factors that have been mentioned in previous paragraphs. For example, limited resources can force investigators to identify causal factors that are characteristic of a class of incidents rather than analyse an incident for any distinguishing characteristics. Alternatively, organisational pressures can persuade investigators that an incident supports some more general political argument. The ambiguous nature of many incidents can make it difficult to resist such influences. As we have seen, adverse occurrences typically have many potential causes. Given sparse data, limited resources and the pressure to act, it is hardly surprising that some investigators are tempted to 'cut corners'. Such practices often only come to light in the aftermath of a major accident. This is illustrated by the treatment of Signals Passed at Danger (SPADs) on UK railways. Chapter 7 quoted the report from Her Majesty's Railway Inspectorate, which found that "in some cases greater emphasis was placed on completing a multi-page form than getting to the root cause of the SPAD incident" [349]. Incident investigations tended to focus on issues of driver vigilance rather than the placement of signals or on the other protection mechanisms that were intended to prevent these incidents from occurring. The HMRI report concluded, investigators might have looked deeper into these incidents if they had been required to follow more rigorous techniques for root cause analysis.

- it The Problems of Ambiguous and Limited Evidence. Incident reconstructions help to establish *what* happened. Causal analysis then identifies the reasons *why* an incident took place. As we have seen, however, these distinctions are difficult to maintain during an incident investigation. Causal hypotheses are formed and reformed as new evidence is obtained about the course of an incident. This creates problems because the resource limited nature of many enquiries can force investigators to develop ad hoc stopping rules. These involve procedures to help them decide when to stop gathering more evidence in support of their analysis. Typically, these procedures involve team presentations or discussions with safety management who must then authorise the end of an investigation. Other circumstances can prematurely curtail a causal analysis. For instance, there may be little direct evidence about the events that led to an incident. Paradoxically, however, NASA's Mars Polar Lander report demonstrates that a lack of evidence does not bring a causal investigation to a premature conclusion [579]. In contrast, it opens up a vast number of possible explanations that must be discounted before reaching a tentative conclusion. In assessing the analytical techniques that will be presented in this chapter, it is therefore important to remember that investigators may have to use them to discount certain hypotheses as well as to support others.

- *The Problems of Intention.* The previous paragraph has argued that causal analyses are complicated by a lack of evidence about the events leading to a failure. This evidence, typically, relates to the observable behaviour of system components. Similar problems are created when analysts lack information about less visible influences on the course of an incident. In particular, it can be difficult to determine the role that human intention plays in an adverse occurrence. Chapter 3 has introduced numerous distinctions between different forms of error and violation. In practice, however, investigators often lack the information that is necessary to distinguish between these different forms [868]. For instance, mistakes stem from an inappropriate intention. It can be difficult for individuals to admit to such intentions in the aftermath of a near miss incident. These problems also affect the interpretation of human behaviour captured on video and audio logs. For instance, individuals have been observed to act in bizarre and pathological ways. They have disregarded operating procedures and violated safety requirements through factors as diverse as boredom, curiosity and a sense of fun [863]. It seems apparent that the advocates of cockpit video recorders significantly underestimate the problems of interpreting human intentions from the behaviour that is captured by these devices. Pedrali's video analysis of optimal and sub-optimal behaviour in commercial test pilots provides ample evidence of this [672]. Later section will describe how ethnographic and work-place studies have been proposed as means of supporting the eventual analysis of such behaviours.

- *Inter-Analyst Reliability.* Many of the problems described in this section stem from a meta-level concern that investigators should be able to replicate any analysis of an incident. This is supported if investigators justify their decision to use a particular technique to support their causal analysis. They should also document any intermediate findings that emerge to support or refute particular conclusions. These requirements enable others to replicate the application of particular analytical techniques. They will not, of course, enable others to directly replicate the results of any causal analysis. Lekberg's work has shown that these results are not simply determined by the choice of an analytical technique [484]. They are also determined by the educational background of the investigator. McElroy has provide a preliminary validation of these ideas [529]. His work showed that even when analysts are trained to use one of the more advanced techniques for causal analysis, their findings will vary considerably even for the same incident. Such problems can be addressed by ensuring that the analysis is replicated by a sufficient number of analysts. This form of mass replication can be used to minimise individual differences in interpretation. However, this averaging out can often lead to polarised views within a team of investigators and it is not clear that a consensus must emerge from replicated forms of analysis. In addition, most reporting systems cannot afford to validate their conclusions through the repeated replication of a causal analysis. There can, therefore, be little confidence that any of the techniques in this chapter will ensure inter-analyst reliability. This is true even for techniques that are supported by formal proof techniques; investigators may disagree about the choice of abstractions that are used within a model. Causal reasoning techniques do, however, increase the transparency of any investigation. They help to document the methods that were used to support particular findings about the causes of an adverse occurrence.

The previous paragraphs provide a stark assessment of the many problems that complicate the causal analysis of safety-critical incidents. These range from pragmatic issues of funding and resource management to the more theoretical barriers to interpreting intentions from observations of human behaviour. Later sections in this chapter, therefore, review some of the solutions that have been proposed to address some of these concerns. In contrast, the following pages describe two incidents that are used to illustrate this comparative study of analytical techniques.

10.1.3 Loss of the Mars Climate Orbiter & Polar Lander

In 1993, NASA commissioned a program to survey the planet Mars. The Jet Propulsion Laboratory (JPL) was identified as the lead centre for these missions. Lockhead Martin Astronautics was selected as the prime contractor. The program initially consisted of the Mars Global Surveyor (MGS), to be launched late in 1996. This global mapping mission is currently orbiting Mars. The Mars

Surveyor'98 project was intended to build on the Global Surveyor's work. This program consisted of the Mars Climate Orbiter and the Mars Polar Lander. Both missions were to satisfy tight financial constraints by exploiting innovative technology under NASA's *faster, better, cheaper* management initiative [570].

The Mars Climate Orbiter was launched in December 1998. It was intended to be the first interplanetary weather satellite. It also had a secondary role to act as a communications relay for the Mars Polar Lander. The Climate Orbiter was to have fired its main engine to achieve an elliptical orbit around Mars in September 1999 [570]. The intention was that it should spend several weeks 'skimming-through' the upper atmosphere. This aero-braking techniques was to achieve a low circular orbit using friction against the spacecraft's solar array to reduce the orbital period from fourteen to two hours. It was during the Mars Orbit Insertion (MOI) maneuver that the Climate Orbiter was lost. The investigation team describe how:

> "During the 9-month journey from Earth to Mars, propulsion maneuvers were periodically performed to remove angular momentum buildup in the on-board reaction wheels (flywheels). These Angular Momentum Desaturation (AMD) events occurred 10-14 times more often than was expected by the operations navigation team. This was because the MCO solar array was asymmetrical relative to the spacecraft body as compared to Mars Global Surveyor (MGS) which had symmetrical solar arrays. This asymmetric effect significantly increased the Sun-induced (solar pressure-induced) momentum buildup on the spacecraft. The increased AMD events coupled with the fact that the angular momentum (impulse) data was in English, rather than metric, units, resulted in small errors being introduced in the trajectory estimate over the course of the 9-month journey. At the time of Mars insertion, the spacecraft trajectory was approximately 170 kilometers lower than planned. As a result, MCO either was destroyed in the atmosphere or re-entered heliocentric space after leaving Mars atmosphere. "[564]

The subsequent inquiry identified twelve recommendations for the development and operation of the Polar Lander. These were addressed by the creation of a Mission Safety and Success Team that drew upon fifty of the Jet Propulsion Laboratory's senior staff. A 'red team' was also created to chart all activities that were intended to feed the lessons of the Climate Orbiter incident into the Polar Lander project.

The Mars Polar Lander was launched approximately three months after the loss of the Climate Orbiter in January, 1999. The same cruise stage was to carry the Polar Lander and two smaller probes that were known as Deep Space 2. This was a highly innovative mission that intended to show that miniaturised components could conduct scientific experiments in space. Deep Space 2 consisted of two micro-probes that were to be released from the Polar Lander before it entered the Mars upper atmosphere. These contained a micro-telecommunications system that was designed to communicate with the orbiting Mars Global Surveyor after the probes had impacted with the planet surface. The Polar Lander and the Deep Space 2 probes approached Mars in December 1999. A final trajectory-correction maneuver, TCM-5, was executed six and a half hours before estimated entry. At 12:02 PST, the spacecraft assumed the its entry attitude. A development decision had previously determined that telemetry data would not be collected during the entry, descent and landing phase. In consequence, the change in attitude had the effect of pointing the antenna away from Earth and the signal was lost, as expected. The Polar Lander was expected to touchdown at 00:14 PST and data transmission was scheduled to begin twenty-four minutes later. Data from the DS2 probes was expected to begin at 07:25 No communications were received from either the Polar Lander or the Deep Space 2 probes. The investigation team reported that:

> "Given the total absence of telemetry data and no response to any of the attempted recovery actions, it was not expected that a probable cause, or causes, of failure could be determined. In fact, the probable cause of the loss of MPL has been traced to premature shutdown of the descent engines, resulting from a vulnerability of the software to transient signals. Owing to the lack of data, other potential failure modes cannot positively be ruled out. Nonetheless, the Board judges there to be little doubt about the probable cause of loss of the mission." [579]

These 'failure' of these two missions provides the case study for the remainder of this chapter. A number of motivating factors help to justify this decision. For instance, these incidents provide a rare insight of the way in which organisations must quickly respond to previous incidents. The Jet Propulsion Laboratory and Lockhead Martin had very limited amounts of time to respond to the loss of the Climate Orbiter before the Polar Lander had to be launched. These examples have, however, been deliberately selected for a number of other reasons. They illustrate the failure of leading-edge technology. Previous chapters have shown that the failure of apparently simple technology can be caused by many complex factors. The Allentown explosion discussed in Chapter 9 provides an instance of this. The gas line did not rely upon particularly complex technology. However, the incident involved regulatory and organisational failure in the decision not to deploy protective devices and warning systems. The explosion also illustrated complex communication problems between the utility supplier, the excavators, the property owners etc. The immediate causes also reflect a failure in communication and training involving the excavation team and the fire inspectors. The complexity of the modelling in the previous chapter reinforces this meta-level point that even simple technology typically has complex failure modes. In contrast, the loss of the Mars missions provides a completely different challenge. These systems were deliberately designed to 'push the technological boundaries' under NASA's *faster, better, cheaper* management initiative [570].

It is important to address a number of objections that can be made to the inclusion of these incidents. Neither of the Mars Surveyor'98 missions resulted in 'near misses'. Both involved significant losses in terms of financial resources and in terms of the opportunity costs associated with their scientific objectives. It is important to emphasise, however, that the principle objective in this chapter is to provide readers with a comparative assessment of different analysis techniques. The focus is, therefore, on the analytical techniques rather than the incidents themselves. The same motivations justified the use of the Allentown explosion to illustrate alternative modelling notations in Chapter 9. The decision to focus on the Mars Climate Orbiter and the Polar Lander is also justified by NASA's publication policy. Readers can access a mass of primary and secondary material. I do not know of any near-miss incident that might provide similar opportunities.

Further objections arise because neither of the Mars Surveyor'98 missions posed a direct threat to human safety once it had left the earth's orbit. It can, therefore, be argued that neither incident is 'safety-critical'. These two case studies can, however illustrate the application of safety-critical techniques to analyse mission-critical failures. The Mars Climate Orbiter and Polar Lander also illustrate how safety-critical techniques can be applied more generally to understand the causes of technological failure. This is not simply a spurious argument about the theoretical value of safety-critical techniques for mission critical applications. It is a pragmatic observation that has been recognised by many industries. The investigation boards that investigated the loss of the Mars missions were governed by the same regulations that cover investigations into the injury and death of civil-service employees and the general public. NASA Procedures and Guidelines document NPG:8621.1 introduced the term 'mishap' to cover these two aspects of mission critical and safety-critical failure [571].

Mission-critical failures provide insights into the possible causes of future safety-critical incidents. This can be seen as a corollary of the previous point. Many analysis techniques reveal common causes of managerial and regulatory failure. As a result, safety and mission-critical incidents may only be distinguished by their consequences rather than by their causes. Leveson reflects this ambiguity when she defines safety to be 'freedom from loss' rather than 'freedom from injury' [486]. The practical consequences of this have again been recognised by many organisations. For instance, one of the principle findings of the Presidential Commission into the the loss of the space shuttle Challenger was that NASA should establish an Office of Safety, Reliability and Quality Assurance [713]. This agency is intended to have direct authority for safety, reliability, and quality assurance throughout the agency and is independent of other NASA program responsibilities. Such initiatives illustrate the perceived importance of integrating safety concerns into wider quality assurance techniques.

There is little published information about the common causes of safety-related and mission-critical incidents. Previous chapters have mentioned Wright's preliminary studies, which suggest that accidents may have different causes than incidents [874]. By extension, it can be argued that safety-related incidents may have different underlying causes that mission-critical failures. In particular, it can be argued that mission critical incidents stem from other aspects of dependability, such as security or availability, that have little to do with safety-related failures. Sadly, more time has been spent on debating the semantics of terms such as 'dependability' than has been spent on determining underlying differences between mission-critical and safety-critical failure. Much of the discussion focuses on the problems of measuring improvements in such as abstract notion when it can be influenced by many more detailed factors including reliability, safety, security, availability etc [475, 486]. For example, a security improvement might increase the dependability of a system in some abstract sense. It can also jeopardise safety if operators are prevented from accessing necessary functions during a systems failure. This debate reflects divisions within the academic community. It also reflects pragmatic distinctions that shape organisational responses to technological failure. For example, NASA's Office of Safety and Mission Assurance provides a common focus for dependability concerns. This organisation does not, however, derive abstract measures of dependability. The focus is on gathering and analysing more detailed information about the causes of mission success and failure. Brevity prevents a more detailed analysis of the practical implications of distinctions between the various components of dependability. In contrast, our focus is on determining whether similar analytical techniques can provide insights into both safety-critical and mission-critical incidents. At present there is insufficient evidence to prove or disprove this hypothesis. The case studies in this chapter can, however, be usefully compared to previous work in incident analysis [409, 470]. Although the analysis presents a single view upon two isolated case studies, there are many strong similarities between the detailed causes of these mission failures and the causes of safety related incidents that were identified in Chapter 3. This should not be surprising given that these safety-related factors are often presented as generic causes of technological and managerial failure.

10.2 Stage 1: Incident Modelling (Revisited)

This section introduces what the US Department of Energy has described as the 'core' analytical techniques for incident and accident investigation. Figure 10.1 provides an overview of these techniques.

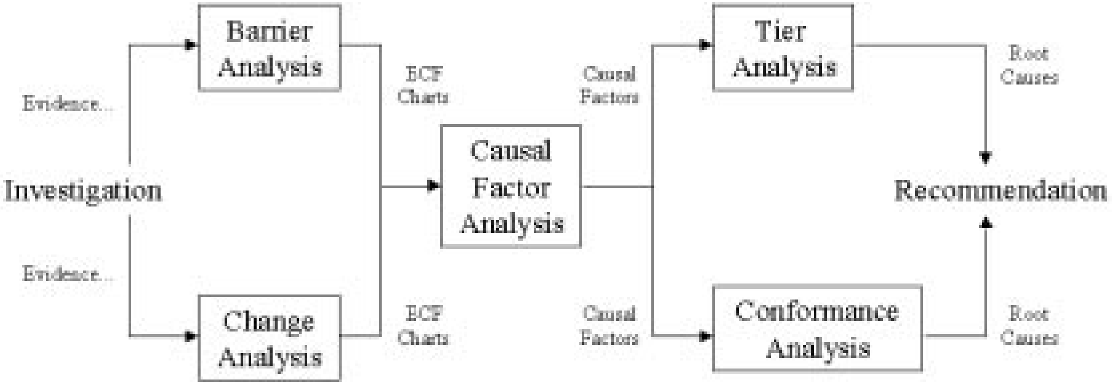

Figure 10.1: Overview of the Dept. of Energy's 'Core' Techniques

The following pages focus on the modelling techniques that form a precursor to any subsequent causal analysis. In order to understand *why* an incident occurred, it is first necessary to determine *what* happened. These are illustrated on the left hand side of Figure 10.1. Unfortunately, the expressive power of these modelling notation is not as great as some of those introduced in Chapter 9. As we shall see, it can be difficult to represent and reason about detailed temporal relationships between the events that are represented in these 'core' modelling techniques. With these caveats in mind, the following sections show how event and causal analysis charts can be used to represent the products of barrier and change analysis. The resulting diagrams then support a more detailed root cause analysis.

10.2.1 Events and Causal Factor Charting

Event and Causal Factor (ECF) charts provide a graphical means of representing the sequence of events leading to a failure. These charts are then annotated with additional causal information. For now, however, it is sufficient to observe that the motivating factors that justify the maintenance of these charts are the same as those for the techniques introduced in Chapter 9:

> "Constructing the events and causal factors chart should begin immediately. However, the initial chart will be only a skeleton of the final product. Many events and conditions will be discovered in a short amount of time, and therefore, the chart should be updated almost daily throughout the investigative data collection phase. Keeping the chart up to date helps ensure that the investigation proceeds smoothly, that gaps in information are identified, and that the investigators have a clear representation of accident chronology for use in evidence collection and witness interviewing." [207]

Figure 10.2 provides a high-level view of the components of an events and causal factor chart. A number of guidelines support the development of these diagrams [207]. The process begins by mapping out a chronology of events. Time is assumed to flow from the left of the diagram to the right. Events represent actions and should be stated with one noun and one active verb. They should be quantified "as much as possible and whenever applicable". The examples suggest that analysts specify how far a worker falls rather than only state that the fall occurred. Times and dates must also be noted and the events should "be derived from" the events that precede them. The approach, therefore, has strong similarities with the use of timelines in previous chapters. Analysts must, however, also distinguish a primary chain from other sequences of events that contribute to the failure. These secondary chains are drawn above the primary line. Without tool support, the problems of maintaining complex graphical structures can limit the scope for introducing these additional event sequences.

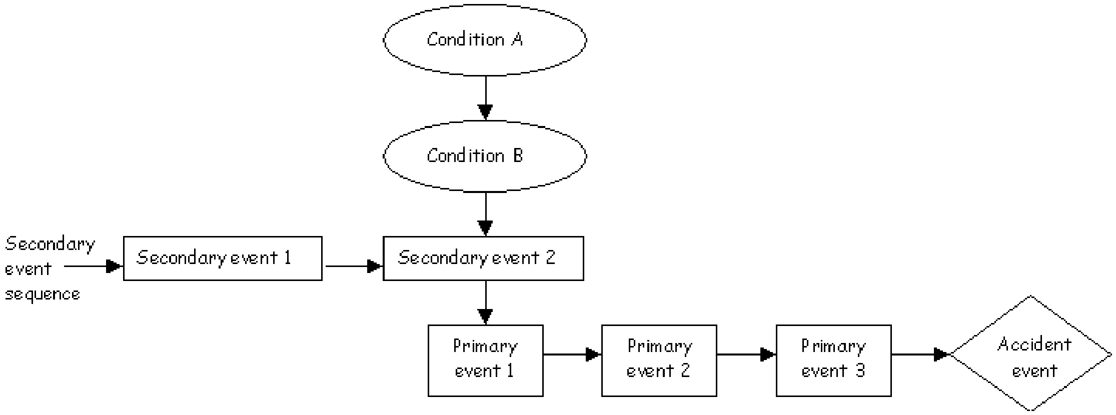

Figure 10.2: Simplified Structure of an ECF Chart

As mentioned, ECF charts have a superficial similarity to timelines. Both exploit linear structures to denote the flow of events leading to an incident or accident. Both approaches must, therefore, consider how to represent state-based information and emergent properties that develop slowly over time. In the case of ECF Charts, these are denoted by the conditions that appear in the ellipses of Figure 10.2. Conditions are passive. For example, they denote that 'there was bad weather' or that 'workers were tired'. They are also associated with the particular events that they help to influence.

Figure 10.3 presents the component symbols that are used in ECF Charts. As with our use of modelling notations, this approach needs to be adapted to support incident analysis. For instance, the diamond used to denote an accident in Figure 10.3 can be used more generally to represent the potential outcome of a 'near miss' incident. Similarly, it is likely that there will be far more

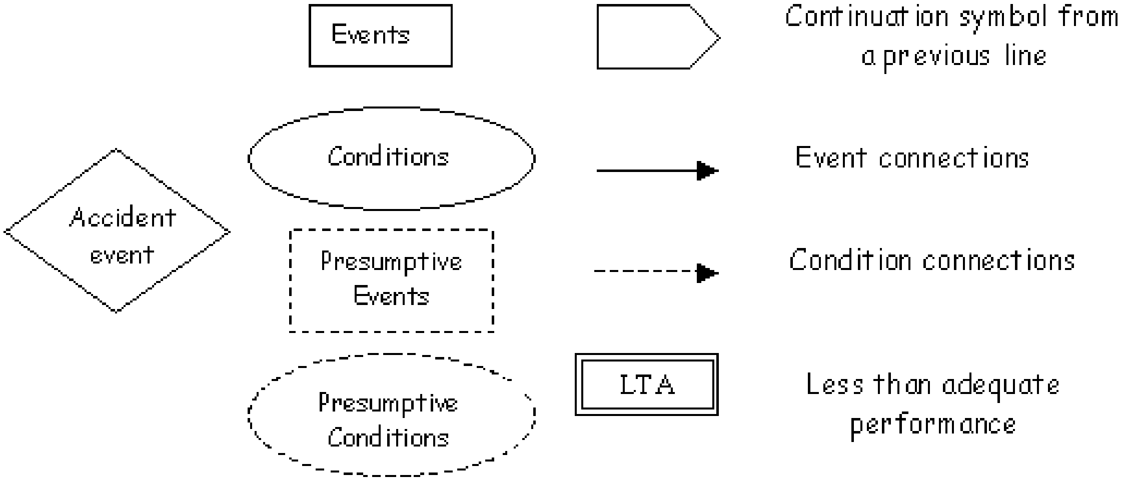

Figure 10.3: Components of ECF Chart

presumptive events and conditions in certain types of incident report systems. For example, analysts are more likely to be forced to make inferences about the events leading to an incident if they have to piece together information from a single submission to an anonymous system. Figure 10.4 illustrates how the ECF notation can be applied to represent the loss of the Mars Climate Orbiter. The intention is to illustrate the information that might be available to investigators in the immediate aftermath of an incident. As can be seen, the primary flow of events is assumed to begin with the launch of the mission on the 11th December. Subsequent analysis will extend the scope of events to consider decisions that were made prior to launch. However, such information may not immediate be available immediately after such an incident. The mission progressed until the last signal was received at 09:04:52.

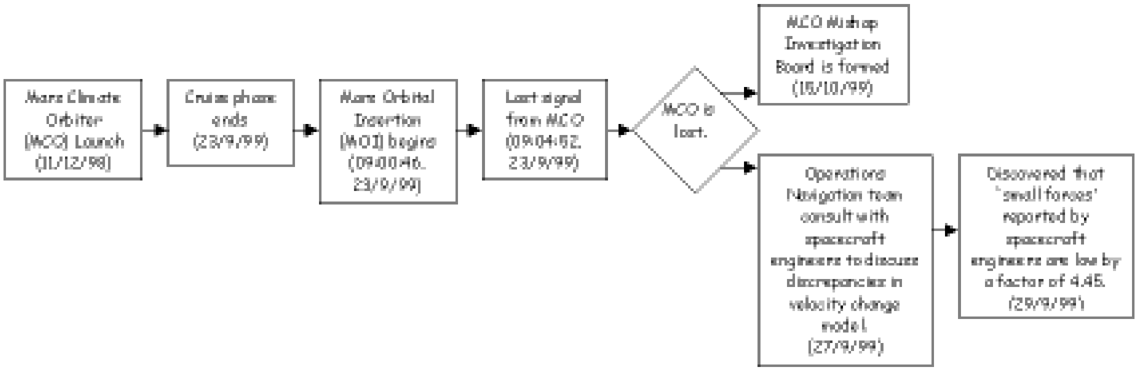
Mars Climate Orbiter (MCO) Launch (11/12/98)
Cruise phase ends (23/9/99)
Mars Orbital Insertion (MOI) begins (09:00:46. 23/9/99)
Last signal from MCO (09:04:52. 23/9/99)
MCO is lost.
MCO Mishap Investigation Board is formed (15/10/99)
Operations Navigation team consult with spacecraft engineers to discuss discrepancies in velocity change model. (27/9/99)
Discovered that 'small forces' reported by spacecraft engineers are low by a factor of 4.45. (29/9/99)

Figure 10.4: High-Level ECF Chart for the Mars Climate Orbiter (MCO)

A number of comments can be made about the use of the ECF notation in Figure 10.4. The accident symbol is used to denote the loss of the Climate Orbiter; MCO is lost. It does not describe the nature of the incident in great detail. NASA investigators considered two possible scenarios; either the craft was destroyed in Mars' atmosphere or it re-entered heliocentric space. These are not shown here because we do not know whether these possible incidents actually took place. This ambiguity stems from NASA's decision not to relay telemetry data during Mars Orbit Insertion. The same decision was taken during the development of the Polar Lander. This deliberate design feature reduced project development costs but clearly also reduced the information that was available to subsequent investigators. As the analysts commented “the decision not to have EDL telemetry was a defensible project decision, but an indefensible programmatic one.” [579].

A second important feature of Figure 10.4 is the way in which it extends beyond the loss of the MCO’s signal. The Operational Navigation team met with Spacecraft Engineers to discuss what might have caused the apparent mission failure. This meeting formed part of an initial response that was intended to devise a way of re-establishing contact with the mission and then, later, to learn any immediate lessons that might affect the Mars Polar Lander. Shortly after this meeting, a bug was discovered in the ‘Small Forces’ software that formed an important component of the navigation system. This sequence of events is critical to any understanding of the MCO incident, not simply because it helped to identify the probable cause of the failure but also because it took place *before* the NASA Mishap investigation board had been formed.

It is inevitable that informal analysis will be conducted in the aftermath of many incidents. In particular, the limited launch window for the Mars Polar Lander made it imperative that lessons were learned as quickly as possible. It can also be argued that by discussing the causes of failure, engineers can make the best use of any opportunities to mitigate the consequences of an incident. However, there also a number of concerns about such interim forms of analysis. Firstly, operators may actually exacerbate the situation if they intervene with partial knowledge about the causes of an incident. The Chernobyl and Three Mile Island accidents provide graphic illustrations of this point. In the former case, Soviet operators exacerbated their problems by rapidly inserting control rods into the reactor that had previously been almost fully withdrawn. Rather than dampening the reaction, positive void coefficients created the opposite effect. Operator intervention at Three Mile Island led the NRC to specify that users should not intervene in similar circumstances without a sufficient period to formulate a detailed diagnosis of the causes of the failure [219]. Secondly, there is a danger that groups who are involved in an incident may prepare an explanation of the failure that cannot be supported by a more detailed analysis. At its most extreme, this may extend to collusion in falsifying evidence. At its most benign, the identification of a probable cause by groups of workers in the aftermath of an incident can have the effect of biasing, or blinkering, any subsequent investigation. Neither of these objections can be applied to the MCO engineers or to NASA’s Mishap Investigation board. It should be noted, however, that the MCO phase I report focuses almost exclusively on the faults identified by the Operational Navigators and the Spacecraft Engineers following their meeting on the 27th September.

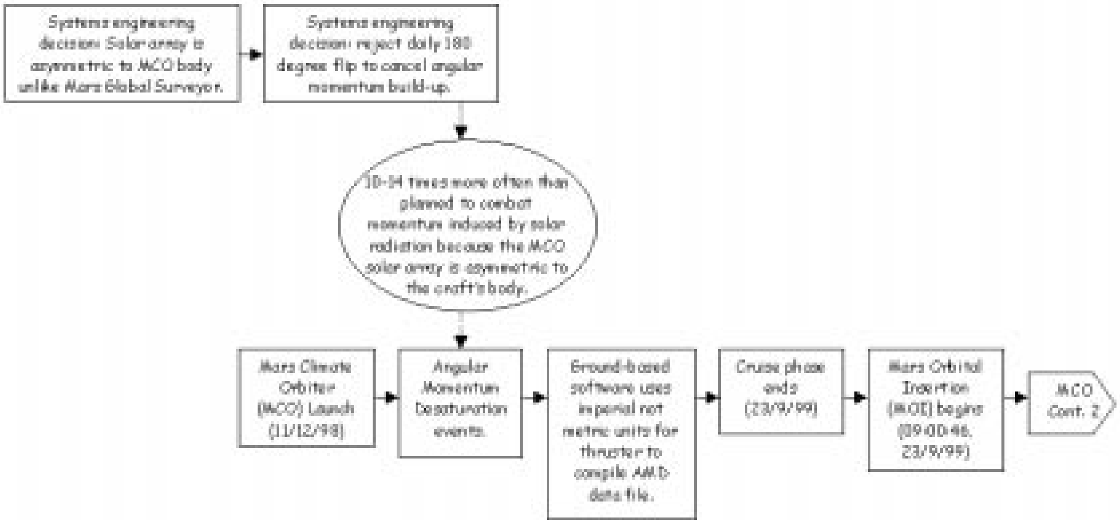
Systems engineering decision: Solar array is asymmetric to MCO body unlike Mars Global Surveyor.
Systems engineering decision: reject daily 180 degree flip to cancel angular momentum build-up.
10-14 times more often than planned to combat momentum induced by solar radiation because the MCO solar array is asymmetric to the craft's body.
Mars Climate Orbiter (MCO) Launch (11/12/98)
Angular Momentum Desaturation events.
Ground-based software uses imperial not metric units for thruster to compile AMD data file.
Cruise phase ends (23/9/99)
Mars Orbital Insertion (MOI) begins (09:00:46, 23/9/99)
MCO Cont. 2

Figure 10.5: Angular Momentum Desaturation Events Affect MCO Navigation

Figure 10.5 extends the previous ECF chart to illustrate an interim stage in the analysis of the MCO incident. As can be seen, this diagram focuses in on events between the launch and the completion of the cruise phase. In particular, it focuses on Angular Momentum Desaturation events. These maneuvers were partially determined by the 'Small Forces' software. As Figure 10.4 shows, this was the code that had been identified as the potential problem by the Operational Navigators and the Spacecraft Engineers. Figure 10.5 shows that ground based software used pounds of force per second rather than Newtons per second to represent thruster performance. This code was used to generate the Angular Momentum Desaturation file that was then used as input to subsequent navigation software and so repeated AMD events would compound any inaccuracies. The condition above the AMD event denotes the observation that Angular Momentum Desaturation maneuvers had to be carried 10 to 14 times more often that had been planned. This was to counter-act the momentum that was induced by radiation acting on the spacecraft's solar array. As can be seen, a secondary line of events explains why AMD maneuvers were so common. A decision was taken to use asymmetric solar panels. this was different to the symmetric configuration used on the Mars Global Surveyor. The frequency of AMD events on the MCO also stemmed from a decision not to perform what were termed 'barbecue' maneuvers in which the craft was flipped through 180 degrees every twenty-four hours.

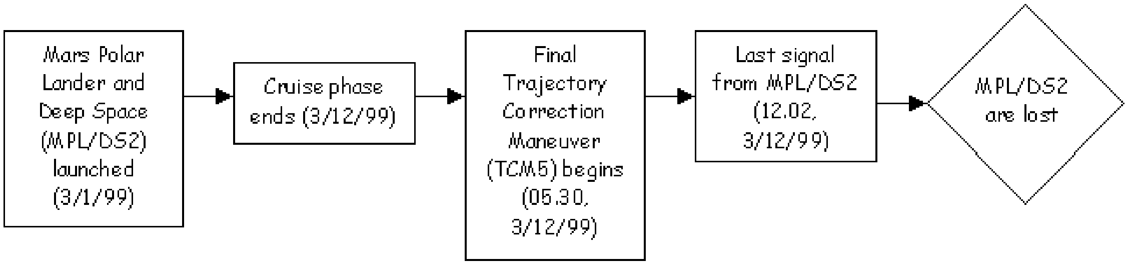
Mars Polar Lander and Deep Space (MPL/DS2) launched (3/1/99)
Cruise phase ends (3/12/99)
Final Trajectory Correction Maneuver (TCM5) begins (05.30, 3/12/99)
Last signal from MPL/DS2 (12.02, 3/12/99)
MPL/DS2 are lost

Figure 10.6: High-Level ECF chart for the Mars Polar Lander (MPL)

Previous ECF charts have focussed on the loss of the MCO. In contrast, Figure 10.6 presents a very high-level view of the observable events that took place before the loss of the Mars Polar Lander. It is important to note again that this diagram does not represent the exact events that might have contributed to the loss of the Lander and the Deep Space 2 probes. The Mars Polar Lander and Deep Space 2 missions might have been destroyed in the atmosphere or re-entered heliocentric space. They might also have been damaged by impact on landing or communications failures might have prevented subsequent communication. The lack of telemetry data can prevent analysts from assessing the likelihood of these different scenarios until a secondary investigation is completed. It is also important to note that this incident is slightly more complex than the loss of the Climate Orbiter. Any failure scenario represented by an ECF chart must account for the loss of the Lander as well as both of the Deep Space 2 mission. Both probes could independently communicate with the Mars Global Surveyor after they had been deployed on the planet surface. A single failure mode is most likely to have occurred prior to the separation of the probes from the Lander. Any failure after separation is most likely to have involved two different failure modes.

Figure 10.7 provides a more detailed view of two of the failure modes that might explain the loss of the Polar Lander and Deep Space 2 missions. As can be seen, the nature and scope of the ECF chart will change as more information becomes available. In this example, the loss of the Polar Lander occurs after the premature shut down of the engines at forty feet from the planet surface. This is influences by a software condition which specified that the engines should be cut if there were two consecutive readings from Hall effect magnetic sensors and the Lander's radar detected that the surface was less than forty meters away. Hall effect sensors were attached to each of the Lander's legs. These were intended to function as follows. Once a leg touched the surface of the planet, the resultant motion would move a magnet away from the sensor. This movement would reduce the

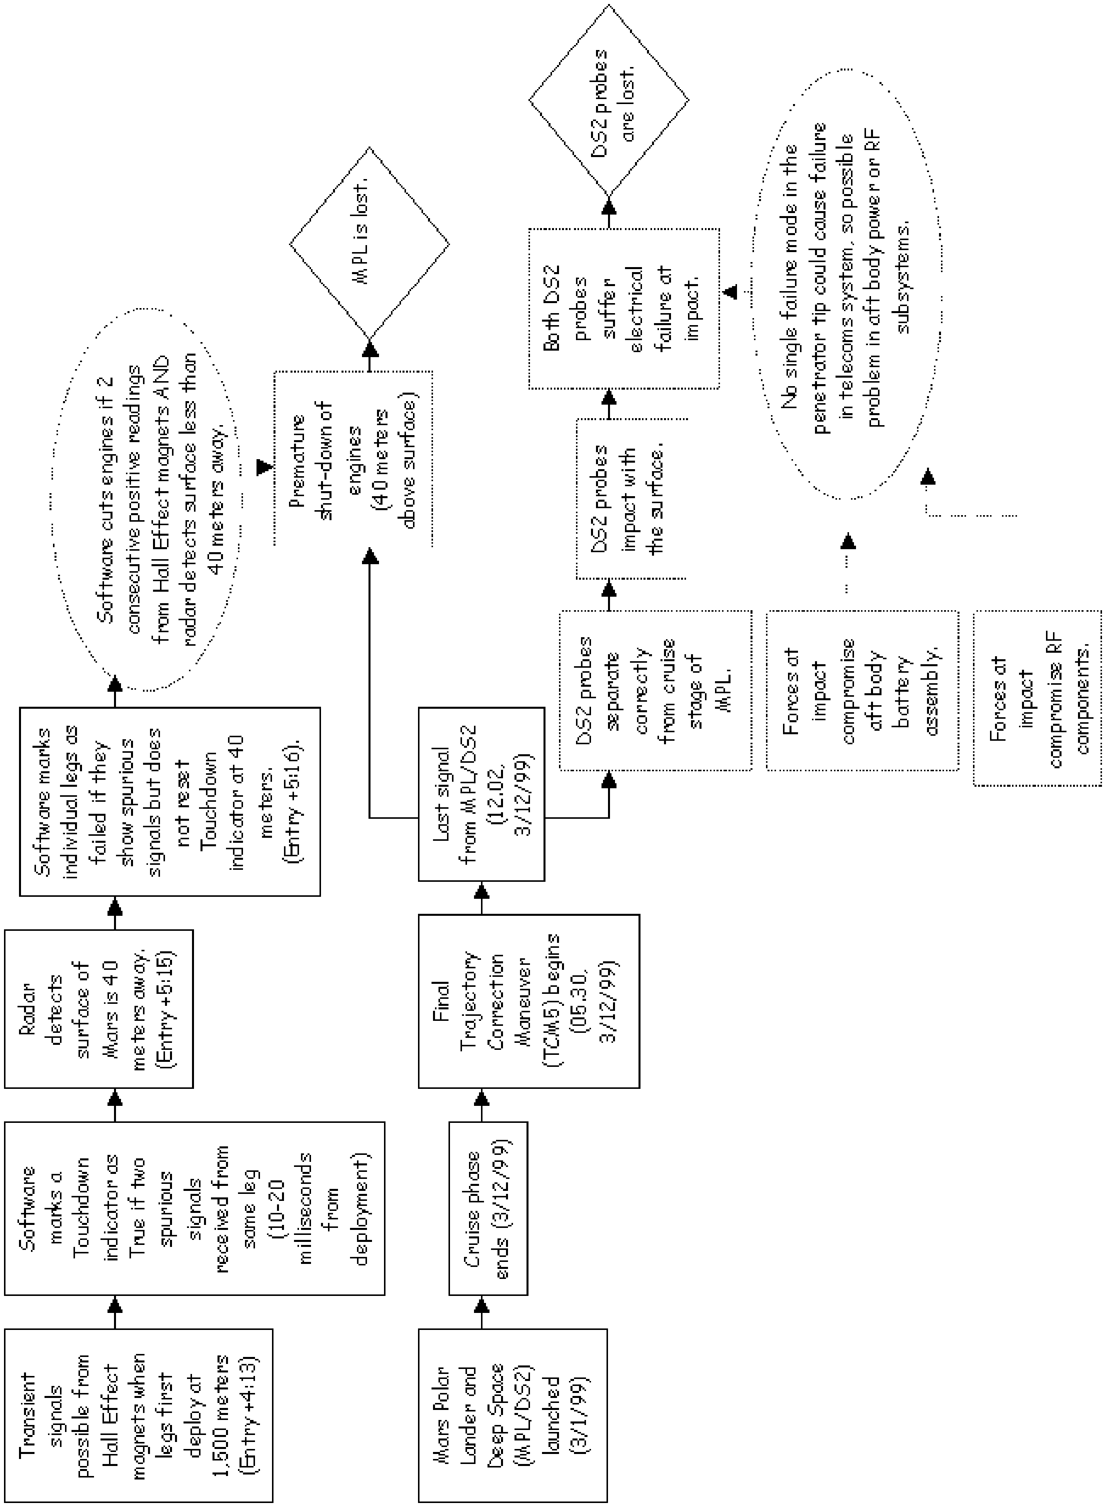
Transient signals possible from Hall Effect magnets when legs first deploy at 1,500 meters (Entry+4:13)
Software marks a Touchdown indicator as True if two spurious signals received from same leg (10-20 milliseconds from deployment)
Radar detects surface of Mars is 40 meters away. (Entry+5:15)
Software marks individual legs as failed if they show spurious signals but does not reset Touchdown indicator at 40 meters. (Entry +5:16).
Software cuts engines if 2 consecutive positive readings from Hall Effect magnets AND radar detects surface less than 40 meters away.
Premature shut-down of engines (40 meters above surface)
MPL is lost.
Mars Polar Lander and Deep Space (MPL/DS2) launched (3/1/99)
Cruise phase ends (3/12/99)
Final Trajectory Correction Maneuver (TCM5) begins (05.30, 3/12/99)
Last signal from MPL/DS2 (12.02, 3/12/99)
DS2 probes separate correctly from cruise stage of MPL.
DS2 probes impact with the surface.
Both DS2 probes suffer electrical failure at impact.
DS2 probes are lost.
Forces at impact compromise aft body battery assembly.
Forces at impact compromise RF components.
No single failure mode in the penetrator tip could cause failure in telecoms system, so possible problem in aft body power or RF subsystems.

Figure 10.7: Premature MPL Engine Shut-Down and DS2 Battery Failure

magnetic field below the sensor's trigger level. However, as can be seen from the upper-left event in Figure 10.7, spurious signals are generated by the sensors when the legs are first deployed into a landing position at some 1,500 meters from the surface. To prevent this from have a disastrous effect, the software systems disregard any signals that are received from the Hall effect sensors until the on-board radar detects the surface at less than forty meters above the surface. The ECF chart in Figure 10.7 represents a possible failure sequence for this approach. If the sensors generate two consecutive spurious signals on leg deployment then a variable Touchdown is initially marked as true. This is not reset to False even though the on-board radar detects that the surface is more that 40 meters away. As a result, when the radar eventually does detect that the surface is 40 meters away the software retains the spurious value of the Touchdown signal that was generated during leg deployment. The two conditions in the software are now satisfied and the engines are cut even though none of the legs are in contact with the surface.

Figure 10.7 also represents different events leading to the loss of the Deep Space 2 probes. These probes would have separated from the Lander long before the engines were cut and so a different explanation has to be found for the loss of any signal between these devices and the Mars Global Surveyor. A presumptive event is used to denote that the probes correctly separated from the Lander. There is no means of being completely sure that this did occur given the lack of telemetry data. A number of alternative failure scenarios can be considered in which the separation did not take place, these would have to be represented in additional ECF chart. In this example, however, correct separation leads to the assumptions that the probes impacted with the planet surface but that both suffered an electrical failure. The associated condition is used to indicate that this is a possible failure scenario because there are no common mode failures in the penetrator section of the probe that could cause a failure in the telecommunications systems. This is a slight simplification if the tethering mechanisms is considered to be part of the penetrator. The loss of both probes can be explained by a failure in either the radio assembly or the battery components that were both located in their aft section.

It is important to stress that the ECF charts in this section provide a very limited view of the possible failure scenarios. In practice, investigators must develop a number of similar diagrams to represent alternative sequences of events. It is important also to remember that the ECF technique was not initially intended to support the analysis of high-technology failures within the aerospace industry. The Polar Lander and Climate Orbiter case studies were deliberately chosen as a challenge to the application of these analytical techniques. For example, the decision not to provide telemetry links during the Lander's Entry, Descent and Landing or the Orbiter's insertion creates a degree of uncertainty that is not often apparent in the more usual application of ECF diagrams to occupational injuries [207].

This section has shown how ECF charts can be used to develop high-level reconstructions of the events that contribute to particular failure scenarios. As can be seen, this involves the identification of observable events, such as the last signals from the Lander, and presumptive events, such as battery damage to the Deep Space 2 probes. These diagrams, therefore, represent an initial stage in the causal analysis of an incident [209]. However, they do not go much beyond the reconstructive modelling techniques that were introduced in Chapter 9. To distinguish between root causes and contributory causes, investigators must recruit a range of complementary analytical techniques. These can be used to ask deeper questions about *why* particular events did or did not contribute to a failure scenario. The results of techniques, such as barrier analysis, can then be used to develop more detailed ECF diagrams.

10.2.2 Barrier Analysis

Barrier analysis has its modern roots in the early 1970's when Haddon proposed a taxonomy of different controls that can be used to mitigate or direct the transfer of energy in safety-critical systems [298]. These included measures to reduce the amount of energy that is generated, measures to separate a target from the source of energy either in time or space, measures to modify shock concentration surfaces and to strengthen the target. These general ideas led to the development of more formal techniques for barrier analysis both as a tool for incident analysis and also as a constructive design tool. As with ECF charting, this technique was driven by the requirements of the US Department of Energy to develop techniques that support the development and analysis of a range of hazardous processes, including nuclear power generation. It is important to stress that barrier analysis also supports the reconstruction and simulation techniques that were described in previous chapters. Fault trees, time-lines, Petri Nets can all be used to capture insights about the successes and failures of potential 'protection devices'. However, barrier analysis is most often used by analysts as a means of extending an initial ECF chart to consider a broader range of potential root causes.

Barrier analysis starts from the assumption that a hazard comes into contact with a target because barriers or controls were unused or inadequate. A hazard is usually thought of as an unwanted energy transfer such as the passage of electricity from an item of equipment to an unprotected worker. Energy can be 'kinetic, biological acoustical, chemical, electrical, mechanical potential, electro-magnetic, thermal or radiation' [207]. The target is the person, equipment or other object that can be harmed by a hazard. Barriers represent the diverse physical and organisational measures that are taken to prevent a target from being affected by a potential hazard. Although distinctions are blurred, many barrier analysis techniques identify controls and safety devices. Control barriers direct wanted or 'desired' energy flows. They include conductors, disconnect switches, pressure vessels and approved work methods. Safety devices are barriers to unwanted energy flows. These include protective equipment, guard rails, safety training and emergency places [208]. The reason that such distinctions can be difficult to make is that the same energy flow might be both wanted and unwanted at different times during an application process. For instance, the Landers thrusters deliver necessary power during the landing sequence. However, this same power source might topple the craft if it continues after the legs have touched the planet surface. The Hall sensors can, therefore, be seen both as controls and safety devices. They acted as a control during the descent because they kept the thrusters working. If the engines were cut then the Lander would be destroyed. However, one the craft has landed the same devices act as safety devices because the power is no longer wanted. Have acknowledged the practical difficulties created by any distinction between safety and control devices, it is possible to distinguish a number of further barriers.

It is possible to identify three different forms of barriers: people; process and technology. For example, material technology has produced physical barriers that directly prevent a hazard from affecting a target. They include guards, gloves and goggles, protective clothing, shields. As we shall see, these devices are often rated to be effective within certain tolerances. For example, a fireguard may provide protection against a fire within particular heat and time limitations. Dynamic barriers include warning devices and alarms [208]. These are not continually apparent but are only issued when the system detects that there may be a potential hazard. This definition can also be extended to include physical interlocks that restrict access or actions during critical phases of an operation. The limitations with this approach stem from the dynamic nature of these warnings. Operators may fail to notice information about a potential hazard. Operators may also choose to disregard or circumvent warnings, especially, if they have been presented with a succession of false alarms. Conversely, warnings may not be invoked even though a hazard may be present. This poses a particular threat if operators grow accustomed to the additional protection afforded by these barriers.

Process barriers include the use of training, of checklists, of standard operating procedures and other forms of workplace regulation that are intended to protect operators and their equipment from potential hazards. Chapter 3 has argued that these procedures can either be explicitly supported by line management or they may arise over time as the result of implicit procedures within everyday working practices. The later class of barriers can be unreliable if new employees fail to observe the way in which existing employees follow these unwritten rules.

People also represent a further class of barrier that can protect a target from a hazard. Human often act as the last barrier against the adverse consequences of energy transfers. The Office of Operating Experience, Analysis and Feedback in the US Department of Energy concludes that:

> "Human action is often, but not always, associated with a procedural barrier. Examples of human action serving to control a hazard are controlling and extinguishing a fire, de-energizing an electrical circuit either in response to a procedure or as part of

> safe work practice, evacuating a building in response to a fire or a criticality alarm, etc." [205].

Managerial and administrative policies can also be interpreted as a form of meta-level barrier. These constraints do not directly protect any particular target from any particular hazard. For instance, they do not directly involve a physical device shielding the operator from a heat source. In contrast, managerial and administrative barriers help to ensure that the acquisition, development, installation and maintenance of a system ensures the adequate provision of more direct barriers to protect potential targets.

The previous paragraphs have mentioned that there are a number of different ways in which barriers can fail. The following list provides a high-level overview of these failure modes:

- *Barrier is impractical - impossible.* There are situations in which it is impossible to provide adequate barriers against a potential energy transfer. Ideally, such situations are identified during a safety analysis. If the hazard could not be prevented or mitigated, regulators should ensure that the process fails to gain necessary permissions. Payne provides numerous examples of this in his analysis of planning applications for safety-critical production processes [669]. He cites a series of incidents in which it was impossible to protect the public once chemicals had been released into the environment. In retrospect, permission should not have been granted for the processes to be sited within urban developments.

- *Barrier is impractical - uneconomic.* In other circumstances, it may be technically feasible to develop appropriate barriers but their cost may prevent them from being deployed. As we have seen, a spate of 'near misses' and accidents persuaded regulators to back the introduction of a Train Protection Warning System on UK railways. This is estimated to cost approximately £310 million. The more sophisticated Advanced Train Protection system was rejected as being uneconomic, at an estimated cost of £2 billion [690]. The obvious weakness with this form of analysis is that the perceived benefits that are associated with particular barriers can change in response to public anxiety over particular incidents. The Southall and Paddington crashes led to a detailed reassessment of the economic arguments against the introduction of the more advanced system.

- *Barrier fails - partially.* A barrier that has been successfully introduced into an application process may, however, fail to fully protect the target from a potential hazard. This is an important class of failure in many incident reporting systems because it represents situations in which barriers provide some protection but may not, under other circumstances, have prevented the hazard from being relaised. For instance, the Mishap Investigation Board into the loss of the Climate Orbiter directed the Polar Lander team to introduce a series of protective barriers. These included the establishment of a 'red team' that was intended to:

  > "study mission scenarios, to ensure operational readiness and to validate risks... This team provides an independent, aggressive, almost adversarial yet helpful role, addressing all levels of the project from high-level requirements down through subsystem design. Key review items include: ensuring system success and reliability; reviewing overall system design and design decisions; reviewing system safety and reliability analyses and risk assessments; reviewing planned and completed testing; and reviewing operational processes, procedures and team preparation. Red team review results and recommendations are reported to the project manager and the project team, as well as senior level management at the centers." [570]

  While this device undoubtedly helped to protect the Polar Lander against a number of potential hazards, it failed to provide total protection against the failure modes that were identified in the aftermath of this second incident.

- *Barrier fails - totally.* The distinction between partial and total protection depends upon the nature of the application. This can be illustrated by assuming for a moment that the failure scenario in Figure 10.7 is an accurate representation of the events leading to the loss of the

Polar Lander. The on-board systems prevented it from immediately cutting its engines when the Hall effect sensors first detected spurious readings. From this perspective, the software provided partial protection. However, the software completely failed in terms of the overall mission objectives. The protection was insufficient to ensure the safe landing of the craft. This example illustrates how the success or failure of a barrier must be interpreted with respect to the overall safety objectives of the system as a whole. The craft was lost and hence the protection is interpreted to have failed in its intended function.

- *Barrier is not used - not provided.* This describes a situation in which a barrier might have protected a target had it been available. At a prosaic level, the bug in the Polar Lander software could have been removed by the addition of a statement, (IndicatorState = False), when the radar detects the forty meter threshold. This need not have provided total protection for the mission. There are a number of alternative failure modes. For instance, the Lander may have encountered terrain with a slope steep enough to destabilize the craft on landing.

- *Barrier is not used - by error.* Barriers may not be used during an incident even though they are available and might prevent a target from being exposed to a hazard. For example, the Climate Orbiter had a contingency maneuver plan in place to execute a Trajectory Correction Maneuver (TCM5). This was intended to raise the the orbit, in fact the second pariapsis passage, to a safe altitude [570]. TCM5 could have been used shortly before Mars Orbit Insertion as an emergency maneuver. It was discussed verbally before the MOI but was never executed. The NASA investigators commented that "the analysis, tests and procedures to commit to a TCM5 in the event of a safety issue were not completed, nor attempted" [570]. In consequence, the operations team were not prepared for such a maneuver.

The previous paragraphs have introduced a number of high-level concepts: barriers; targets and hazards. We have also identified ways in barriers may fail to protect a target or may not be available to mitigate or control a potential hazard. We have not, however, provided a mechanism by which these general observations can support the causal analysis of adverse occurrences. Nor have we shown how the findings of such an analysis can be integrated into the ECF charts that were developed in the previous section. Barrier tables, such as that shown in Table 10.1, can be used to address this omission.

Hazard: Impact/Re-Entry	Target: Mars Climate Orbiter
Barrier	Reason for failure?
People	Lack of staff
	Changes in management
	Inadequate training/skills
	Poor communication
Process	Separation of development and operations teams
	No systematic hazard analysis
	Inadequate testing
	Lack of oversight
Technology	Incorrect trajectory modelling
	Tracking problems
	Rejection of barbecue mode
	Rejection of TCM-5

Table 10.1: Level 1 Barrier Table for the Loss of the Climate Orbiter.

Table 10.1 provides a high level view of the barriers that were intended to prevent the Climate Orbiter from re-entering heliocentric space or impacting the planet surface. As can be seen, the people, process and technology distinctions are retained from the previous paragraphs. This reflects the key components for *Mission Success First* that was advocated by the NASA mishap investigators. They argued that “every individual on the program/project team (must) continuously employ solid engineering and scientific discipline, take personal ownership for their project development efforts and continuously manage risk in order to design, develop and deliver robust systems capable of supporting all mission scenarios” [570]. Table 10.1 records some of the reasons why the individuals involved in the Climate Orbiter project failed to adequately protect against the potential loss of the mission.

People Barriers

Firstly, there were insufficient staff. The primary investigation found that the staffing of the operations navigation team was less than adequate. In particular, the Mars Surveyor Operations Project was responsible for running the Global Surveyor and the Polar Lander in addition to the Climate Orbiter. The investigation revealed that these divided responsibilities tended to ‘dilute’ the focus on any single mission. This loading had a particular effect on the Climate Orbiter’s navigation team. The two individuals who led this group found it very difficult to provide the twenty-four hour a day coverage that was recommended during critical phases of a mission, such as the Climate Orbiter’s MOI [564]. The loss of the Climate Orbiter led to an increase in the number of navigators who were assigned to the Polar Lander project. In terms of the earlier mission, however, this lack of personnel may have prevented the navigation team from sustaining their investigation into the anomalies that they found between the ground-based and on-board navigation systems. This, in turn, reduced the navigation team’s ability to operate as an effective barrier to any navigational problems that might ultimately threaten the success of the mission.

Barrier analysis can also be used to identify further ways in which individuals failed to prevent the loss of the Climate Orbiter. In particular, changes in management prevented an effective response to the navigation problems. During the months leading up to MOI, the investigators found that the Mars Surveyor operations team had “some key personnel vacancies and a change in top management” [570]. A number of further problems reduced management effectiveness in combating particular hazards. For example, there was a perceived ‘lack of ownership’ by some operations personnel who felt that the mission had simply been passed onto them by the development teams. A key management failure in this process was that the operations team had no systems engineering or mission assurance personnel who might have monitored the implementation of the process. This, in turn, might have helped to improve communication between these different phases of the mission. Poor communication appears as a separate explanation for the way in which human barriers failed to prevent mission failure. The investigators concluded that “the spacecraft operations team did not understand the concerns of the operations navigation team” [564].The operations navigation team appeared to be isolated from the development team and from their colleagues in other areas of operations. Other problems stemmed from the nature of group communications during the cruise phase. For example, the navigation team relied on email to coordinate their response once the conflicts were identified in the navigation data. The investigators were concerned that this use of technology enabled some of the problems to ‘slip through the cracks’.

Primary and secondary investigations also identified inadequate training as a potential reason why staff failed to identify the potential hazard to the mission. This was connected to the lack of key personnel because there was no adequate means of ensuring that new team members acquired necessary operational skills. In particular, there was no explicit mentoring system [570]. The investigators argued that the “failure to use metric units in the coding of the Small Forces ground software used in trajectory modeling...might have been uncovered with proper training” [564]. Such comments are significant because they come very close to the counterfactual arguments that have been associated with root cause analysis [25]. One particularly important area for concern was that the the operations navigation team was not familiar with the attitude control system on-board the Climate orbiter; “these functions and their ramifications for Mars Climate Orbiter navigation were fully understood by neither the operations navigation team nor the spacecraft team, due to inexperience and miscommunication” [570]. This lack of familiarity with spacecraft characteristics had considerable consequences throughout the incident. In particular, it may have prevented the operational navigation team from appreciating the full significance of the discrepancies that were identified.

Table 10.1 summarises the reasons why individuals failed to protect the Climate Orbiter from mission failure. The previous paragraphs have built upon this analysis to explain why lack of staff, changes in management, inadequate training and poor communication had an adverse effect upon potential barriers. We have not shown how the results of this analysis might be used to inform the development of Effects and Causal Factor diagrams. The first problem in incorporating these additional insights is that many of the barriers, described above, relate to distal factors. They influence several of the events in Figures 10.4 and 10.5. A second issue is that barrier analysis, typically, helps to identify additional events that ought to be introduced into an Effects and Causal Factor diagram. This is particularly important because primary investigations often focus on catalytic events rather than events that weakened particular barriers.

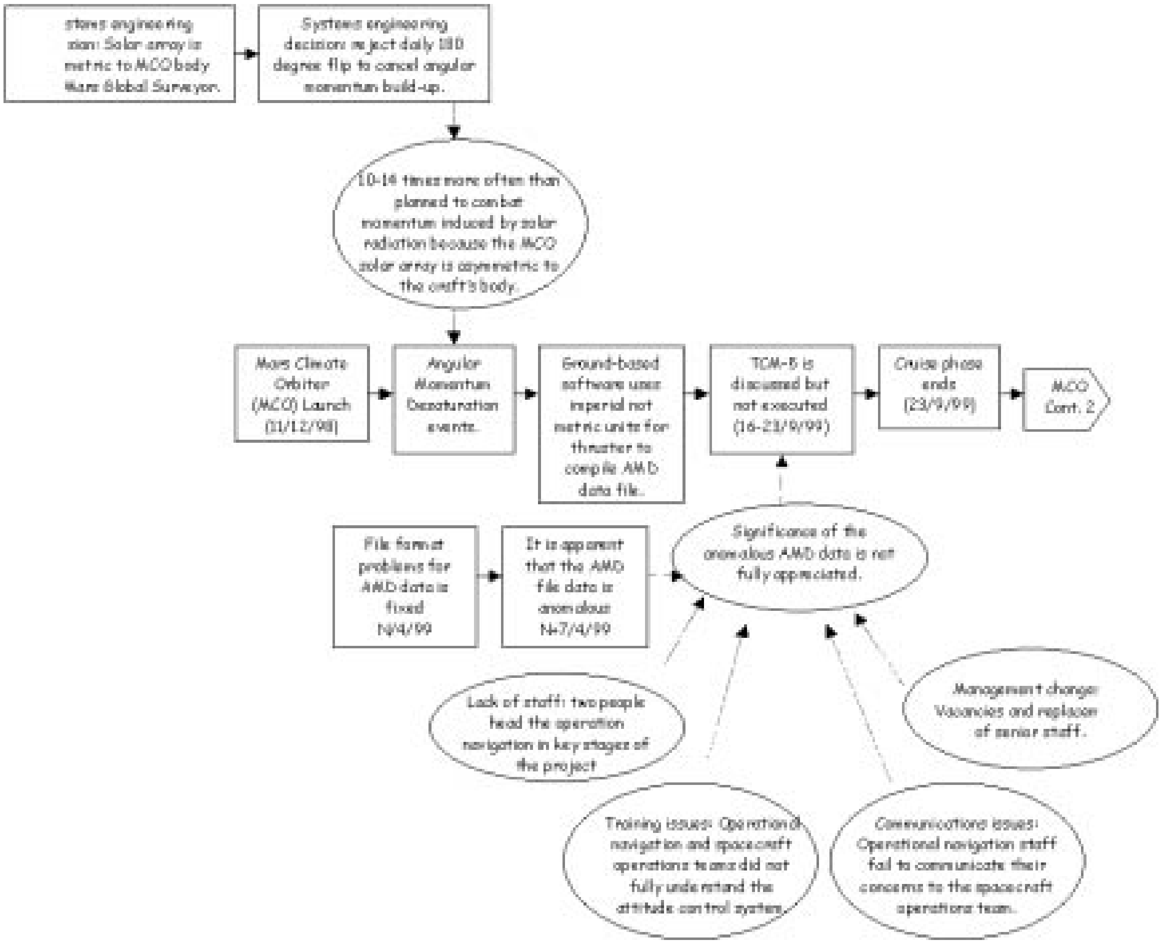
stems engineering sion: Solar array is metric to MCO body Mars Global Surveyor.
Systems engineering decision: reject daily 180 degree flip to cancel angular momentum build-up.
10-14 times more often than planned to combat momentum induced by solar radiation because the MCO solar array is asymmetric to the craft's body.
Mars Climate Orbiter (MCO) Launch (11/12/98)
Angular Momentum Desaturation events.
Ground-based software uses imperial not metric units for thruster to compile AMD data file.
TCM-5 is discussed but not executed (16-23/9/99)
Cruise phase ends (23/9/99)
MCO Cont. 2
File format problems for AMD data is fixed N/4/99
It is apparent that the AMD file data is anomalous N+7/4/99
Significance of the anomalous AMD data is not fully appreciated.
Lack of staff: two people head the operation navigation in key stages of the project
Management change: Vacancies and replacem of senior staff.
Training issues: Operational navigation and spacecraft operations teams did not fully understand the attitude control system.
Communications issues: Operational navigation staff fail to communicate their concerns to the spacecraft operations team.

Figure 10.8: Integrating the Products of Barrier Analysis into ECF Charts

Figure 10.8 integrates our analysis of the human barriers to mission failure into an ECF chart. As can be seen, this diagram introduces a new event into the primary sequence. This denotes the decision not to initiate the TCM-5 maneuver. It was introduced because the previous barrier analysis identified TCM-5 as an important opportunity for preventing the hazard from affecting the target. Figure 10.8 also uses the insights from the barrier analysis to explain why this opportunity was not acted upon. Lack of staff, inadequate training, management changes and poor communication between the operational navigation and spacecraft teams were all factors in the failure to perceive the significance of the AMD data anomaly. Figure 10.8 also illustrates the way in which barrier analysis helps to identify key event sequences that may not have been identified during the initial analysis of an adverse occurrence. As can be seen, this ECF chart has been extended to represent the fact that file formatting errors prevented the navigation team from identifying the AMD anomaly until more than four months after launch.

Process Barriers

Table 10.1 identified four ways in which process barriers may have failed during the Climate Orbiter incident. These related to the separation of the development and operations teams, to the lack of any systematic hazard analysis, to inadequate testing and to the lack of management oversight during particular phases of the mission.

The previous section identified that many of the operational staff lacked necessary training about the operating characteristics of the Climate Orbiter. One reason for this was that the overall project plan did not provide for a careful hand-over from the development project to the operations staff. The Climate Orbiter was also the first mission to be supported by a multi-mission Mars Surveyor Operations Project. The operations staff had to assume control of the Climate Orbiter project without losing track of the Global Orbiter and the Polar Lander missions. These logistical problems were compounded by that fact that the Climate Orbiter project was the first Jet Propulsion Laboratory mission in which only a small number of development staff were 'transitioned' into the operations team. No navigation personnel, made this move from the development of the Climate Orbiter into its operation. This had a number of important consequences for subsequent events during the incident. In particular, the navigation team and other operational staff may have made a number of incorrect assumptions about hardware and software similarities between the Global Surveyor and the Climate Orbiter. The investigators argued that:

> "This apparently caused the operations navigation team to acquire insufficient technical knowledge of the spacecraft, its operation, and its potential impact to navigation computations. The operations navigation team did not know until long after launch that the spacecraft routinely calculated, and transmitted to Earth, velocity change data for the angular momentum desaturation events. An early comparison of these spacecraft-generated data with the tracking data might have uncovered the units problem that ultimately led to the loss of the spacecraft. " [564].

The key point here is that the decision not to transition key development staff into the operation phase removed one of the procedural barriers that otherwise protect JPL missions. The navigational operations team might have realised the potential significance of the AMS anomaly if they had known more about the decisions that had informed the development of the Climate Orbiter.

Figure 10.9 shows how barrier analysis helps to identify a number of additional events and conditions that influenced the course of the incident. The ECF chart has been extended to explicitly denote that a minimal number of development staff were transferred to the operations teams. A number of associated conditions show that the plans for this transition were less than adequate and that this was the first project for the multi-mission Mars Survey Operations project. The previous barrier analysis, however, also raises a number of important questions about the construction of ECF charts. For example, the decision only to transfer a minimal number of staff helped to create the conditions in which operational teams made inappropriate assumptions about the similarity between the Global Surveyor and the Climate Orbiter. These erroneous nature of these suppositions is underlined by the changes in the solar array that are also noted on Figure 10.9. Problems arise because although these incorrect assumptions stem from early in the transition from development to operations, they continue to have an influence throughout the incident. This is difficult to denote use the ECF format introduced in previous section. The condition that represents the potential for incorrect assumptions is surrounded by a double line. Later sections will explain how such conditions provide an important starting point for any subsequent attempts to distinguish root causes from contributory factors.

The hand-over from development to operation was one of several process issues that undermined the Climate Orbiter mission. The lack of any systematic hazard assessment, for instance using Fault

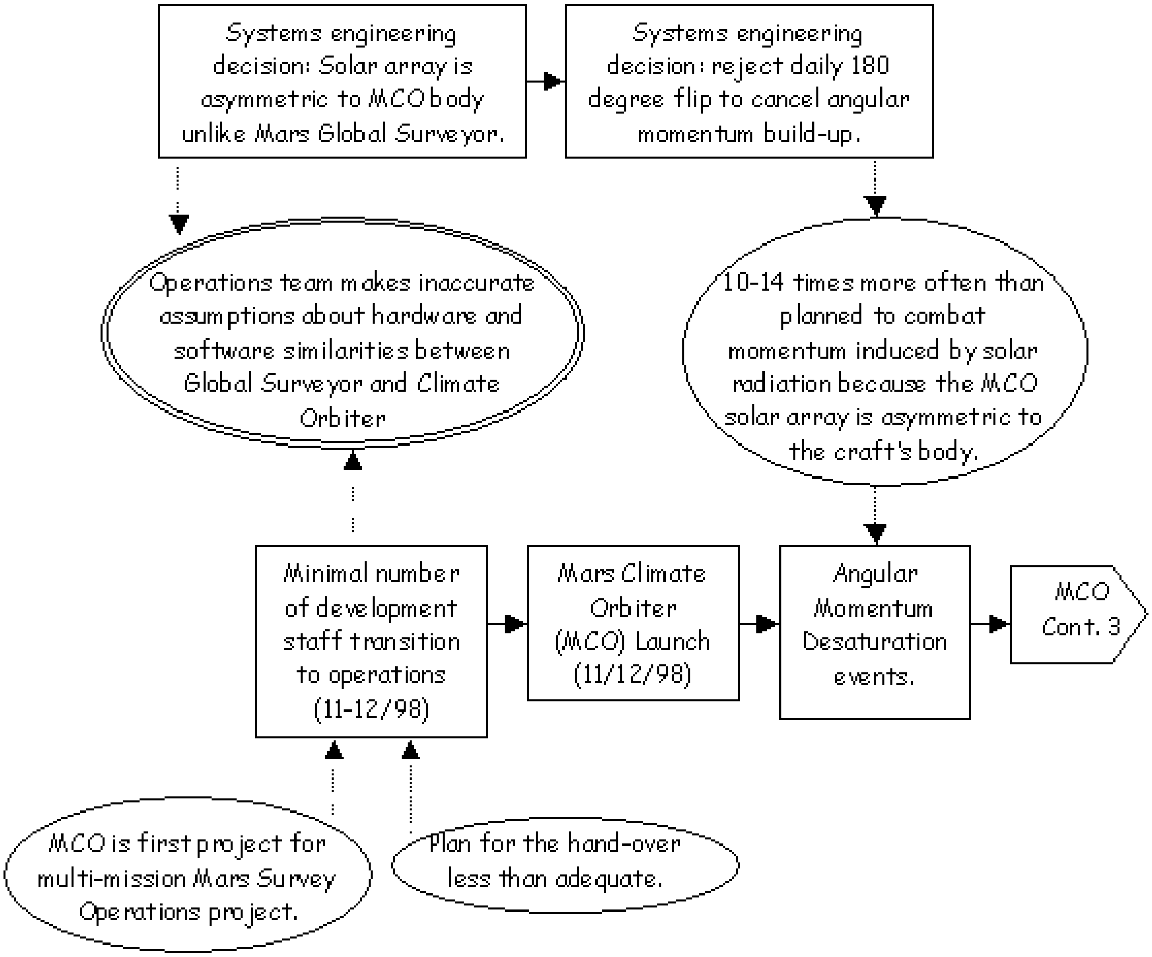
Systems engineering decision: Solar array is asymmetric to MCO body unlike Mars Global Surveyor.
Systems engineering decision: reject daily 180 degree flip to cancel angular momentum build-up.
Operations team makes inaccurate assumptions about hardware and software similarities between Global Surveyor and Climate Orbiter
10-14 times more often than planned to combat momentum induced by solar radiation because the MCO solar array is asymmetric to the craft's body.
Minimal number of development staff transition to operations (11-12/98)
Mars Climate Orbiter (MCO) Launch (11/12/98)
Angular Momentum Desaturation events.
MCO Cont. 3
MCO is first project for multi-mission Mars Survey Operations project.
Plan for the hand-over less than adequate.

Figure 10.9: Process Barriers Fail to Protect the Climate Orbiter

Tree analysis, had numerous consequences for the mission as a whole. This prevented engineers from considering a range of possible failure modes. It also prevented the development and operations teams from conducting a systematic assessment of what were, and what were not, mission critical features. In particular, some form of hazard analysis might have helped to identify that specific elements of the ground software could be 'mission critical' for the operations navigation team. Finally, the lack of a coherent hazard analysis may also have led to inadequate contingency planning. This is particularly apparent in the lack of preparation for TCM-5, mentioned in previous paragraphs. As can be seen, the failure to conduct such an analysis had the knock-on effect of removing a number of potential barriers that might have either detected the navigation software as a critical component prior to launch or might, subsequently, have encouraged operations to reconsider contingency plans once the anomaly had been discovered.

The previous paragraph argued that the lack of any systematic hazard analysis illustrates a further failure of process barriers. Figure 10.10 builds on this analysis by integrating it into the previous ECF charts. This illustrates one of the issues that can complicate the construction of such diagrams. It can be difficult to decide whether or not a particular failure should be represented by the event that triggered the failure or by the conditions that form the consequences of that event. For example, Figure 10.10 include an event labelled Decision not to perform an a priori analysis of what could go wrong on the MCO. This might have been represented by a condition labelled there was no systematic hazard analysis. The ECF manuals provide little guidance on this issue [209, 207]. It is important, however, that some heuristic be used to guide the construction of these

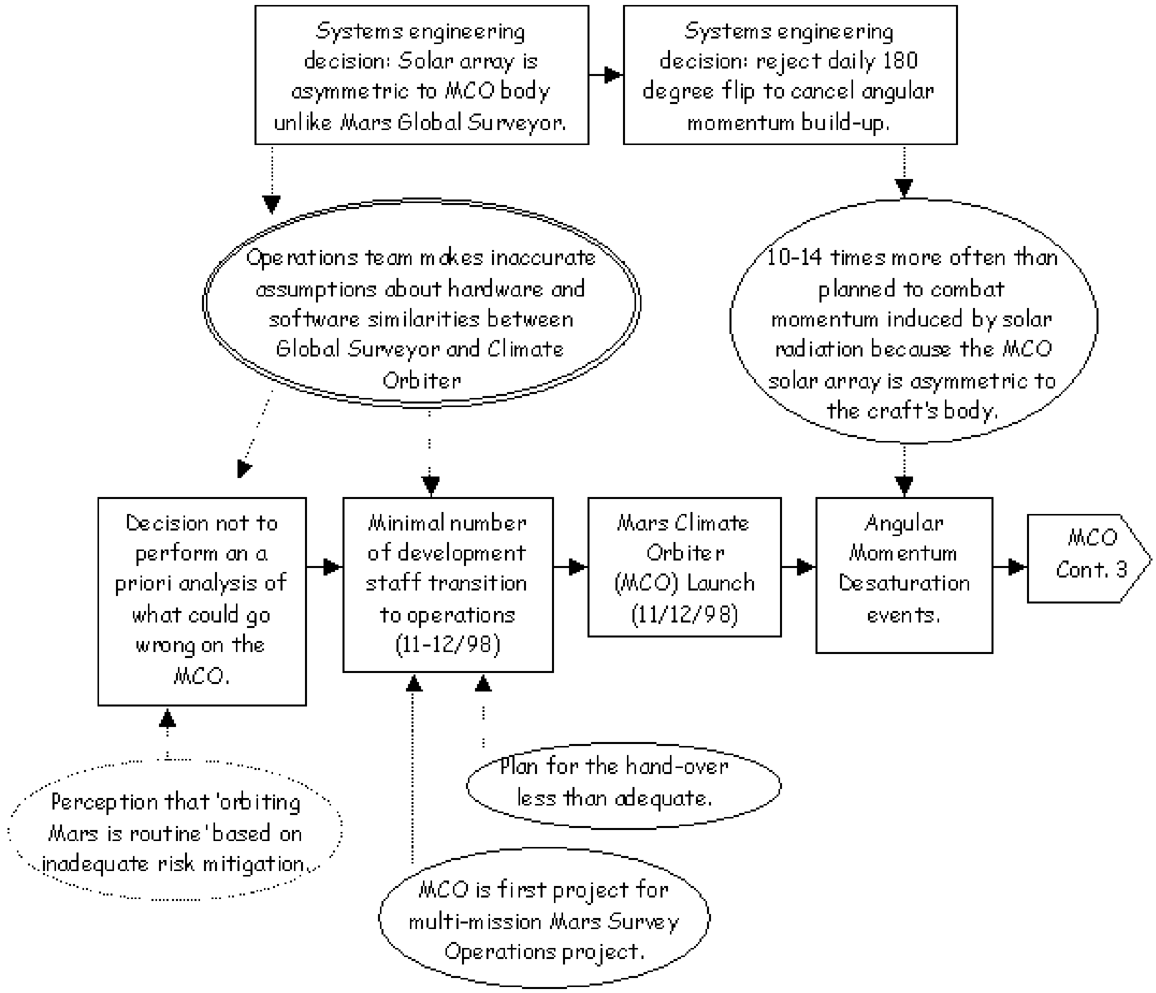
Systems engineering decision: Solar array is asymmetric to MCO body unlike Mars Global Surveyor.
Systems engineering decision: reject daily 180 degree flip to cancel angular momentum build-up.
Operations team makes inaccurate assumptions about hardware and software similarities between Global Surveyor and Climate Orbiter
10-14 times more often than planned to combat momentum induced by solar radiation because the MCO solar array is asymmetric to the craft's body.
Decision not to perform an a priori analysis of what could go wrong on the MCO.
Minimal number of development staff transition to operations (11-12/98)
Mars Climate Orbiter (MCO) Launch (11/12/98)
Angular Momentum Desaturation events.
MCO Cont. 3
Perception that 'orbiting Mars is routine' based on inadequate risk mitigation.
Plan for the hand-over less than adequate.
MCO is first project for multi-mission Mars Survey Operations project.

Figure 10.10: Process Barriers Fail to Protect the Climate Orbiter (2)

diagrams. We have, therefore, use events to denote those stages in an incident that might become a focus for subsequent analysis. Investigators might decide that more needs to be known about the circumstances that influenced any decision not to conduct a systemic hazard analysis. This decision is, therefore, represented as an event rather than a condition.

Further process barriers were undermined by the lack of any sustained validation at a systems level. Navigation requirements were set at too high a management level. In consequence, programmers and engineers were left to determine how best to satisfy those requirements without detailed guidance from others involved in the development process. These problems might not have been so severe had their consequences been detected by an adequate validation process. Several significant system and subsystem flaws were, however, only uncovered after the Climate Orbiter had been launched. For instance, file format errors prevented the navigation team from receiving and interpreting telemetry from the ground system for almost six months. The NASA investigators argued that there was “inadequate independent verification and validation of Mars Climate Orbiter ground software (end-to-end testing to validate the small forces ground software performance and its applicability to the software interface specification did not appear to be accomplished)” [570].

The validation issues and the lack of any system level hazard analysis were exacerbated by a more general lack of oversight during the Climate Orbiter mission. There was little Jet Propulsion Laboratory oversight of Lockheed Martin Astronautics subsystem developments. This created problems as the level of staffing was reduced during the transition from development to operations. Several mission critical functions, including navigation and software validation, received insufficient management oversight. It also became difficult to maintain lines of responsibility and accountability during the project. This point can be illustrated by the Mishap board's description of the relationship between JPL and the contractor:

> "Lockheed Martin Astronautics of Denver, Colorado was selected as the prime contractor. Lockheed Martin Astronautics contracted development responsibilities were to design and develop both spacecraft, lead flight system integration and test, and support launch operations. JPL retained responsibilities for overall project management, spacecraft and instrument development management, project system engineering, mission design, navigation design, mission operation system development, ground data system development, and mission assurance. The Mars Surveyor Project'98 assigned the responsibility for mission operations systems/ground data systems development to the Mars Surveyor Operations Project, Lockheed Martin Astronautics provided support to Mars Surveyor Operations Project for mission operations systems/ground data systems development tasks related to spacecraft test and operations." [564]

Recurring questions in the NASA investigation included 'Who is in charge?' and 'Who is the mission manager?'. The investigators reported repeated examples of 'hesitancy and wavering' whenever individuals attempted to answer the latter question. This is not surprising given the comments made about the feelings of guilt and blame that often operators' reactions to adverse occurrences, see Chapter 5. However, the NASA board also describe how one interviewee answered that the flight operations manager was acting like a mission manager without being designated as such.

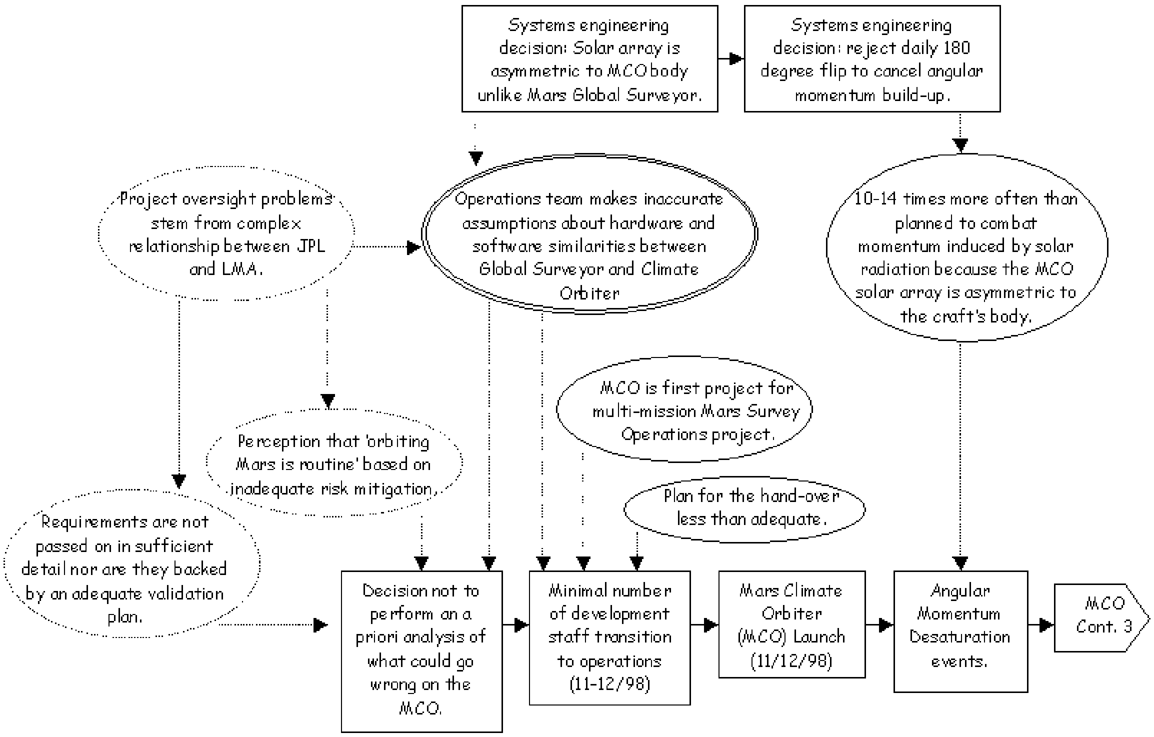
Systems engineering decision: Solar array is asymmetric to MCO body unlike Mars Global Surveyor.
Systems engineering decision: reject daily 180 degree flip to cancel angular momentum build-up.
Project oversight problems stem from complex relationship between JPL and LMA.
Operations team makes inaccurate assumptions about hardware and software similarities between Global Surveyor and Climate Orbiter
10-14 times more often than planned to combat momentum induced by solar radiation because the MCO solar array is asymmetric to the craft's body.
MCO is first project for multi-mission Mars Survey Operations project.
Perception that 'orbiting Mars is routine' based on inadequate risk mitigation.
Plan for the hand-over less than adequate.
Requirements are not passed on in sufficient detail nor are they backed by an adequate validation plan.
Decision not to perform an a priori analysis of what could go wrong on the MCO.
Minimal number of development staff transition to operations (11-12/98)
Mars Climate Orbiter (MCO) Launch (11/12/98)
Angular Momentum Desaturation events.
MCO Cont. 3

Figure 10.11: Process Barriers Fail to Protect the Climate Orbiter (3)

Figure 10.11 shows how the insights that can be derived from a barrier analysis of process failures can be represented within the previous ECF charts. As can be seen the lack of oversight had an important effect on many diverse aspects of the Climate Orbiter's development and operation. It this oversight had been in place then it might have persuaded participants to be more circumspect in their assumptions about the Climate Orbiter's hardware and software characteristics. More coherent oversight might also have encouraged a systemic hazard analysis, especially if more attention had been paid to the validation of high-level requirements.

It should be apparent from the preceding paragraphs that there is no automatic means of propagating the findings of a barrier analysis into the graphical representations of an ECF chart. The investigator must determine how best to translate the findings of their analysis into the events and conditions of Figures 10.10 and 10.11. It, therefore, follows that different investigators might derive different event structures from those shown in this chapter. This introduces a number of concerns about the consistency and validity of any analysis. I am unaware of any research having been conducted into these important aspects of the ECF technique. It can, however, be argued that this analytical process is less about the development of a single coherent view than it is about the explicit representation of what might otherise remain implicit assessments about the success or failure of particular barriers.

Technological Barriers

Technological barriers can also be deployed to support the protection that people and processes provide for safety-critical and mission-critical applications. Table 10.1 has identified four ways in which these technological barriers failed to support the Climate Orbiter mission. There were problems with the trajectory modelling that was intended to identify that potential navigation hazards. The tracking systems that were intended to identify failures in the trajectory models also provided contradictory information. The failure of these barriers became increasingly important because of decisions not to exploit some of the technological measures, including the barbecue mode and TCM-5 contingency, that might otherwise have prevented the mishap from occurring.

The barbecue mode involved a plan to 'flip' the spacecraft by 180 degrees every twenty-four hours. This would have reduced the need for AMD events. The rotation of the aircraft would ensure that any momentum induced by the asymmetric solar panels would have been counteracted in the following twenty-four hours. Previous sections have already shown how this decision can be introduced in an ECF chart, for example Figure 10.5. Similarly, Figure 10.8 introduced the decision not to initiate the TCM-5 maneuver into previous ECF charts. This formed part of an analysis into the failure of people-related barriers. Rather than extend the scope of these previous diagrams, this section focuses on the technological problems that removed navigation and tracking safeguards. Subsequent paragraphs go on to perform a more detailed analysis of the software 'bugs' that removed many of the technological barriers to mission failure.

The previous section has described how problems in the validation of mission critical software created a situation in which several systems had to be debugged during the cruise phase of the mission. This created particular problems because these systems provided important barriers against mission failure. In particular, ground software could not be used to perform the anticipated Angular Momentum Desaturation calculations during the first four months of the cruise. Multiple file format errors were compounded by problems with the data types that were used to represent the spacecraft's attitude. As we have seen, the operations navigation team was forced to use email from the contractor to notify them when a desaturation event was occurring. They then attempted to model the impact on the Climate Orbiter's trajectory using timing information and the manufacturer's performance data. It was not until April 1999 that operations staff could begin using the correctly formatted files. It took a further week for the navigation team to diagnose that the files underestimated the trajectory perturbations due to desaturation events.

The file format and content errors removed important barriers that might otherwise have protected the mission. They prevented the operations navigation team from being able to quickly detect and investigate the underlying calculation problems. These problems might not have had severe consequences if other forms of protection had also been available. In particular, the operations navigation team had limited means of tracking and monitoring the consequences of AMD events. It was difficult to observe the total magnitude of the thrust because of the relative geometry of the thrusters used for AMD activities and the Earth-to-spacecraft line of sight. In consequence, the navigation team had to rely upon the spacecraft's Doppler shift to measure the thrust in this plane. These problems were compounded by the fact that the primary component of the thrust was also perpendicular to the spacecrafts flight path. Changes had to be measured with respect to the craft's original velocity along that plane. These measurement problems stemmed from a navigation strategy that depended on the Earth-based, Deep Space Network to track the Mars Climate Orbiter. A number of alternative technologies might have been used. For instance, the Polar Lander mission also recruited a measurement technique known as 'Near Simultaneous Tracking'. These alternatives were not implemented or were not operational when the Climate Orbiters reached the point of Mars Orbital Insertion [570]. It is important to note, however, that even if they had been implemented they may actually have contributed to the existing confusion about navigation data:

> "The use of supplemental tracking data types to enhance or increase the accuracy of the Mars Polar Lander navigation solutions was discussed. One data type listed in the Mars Polar Lander Mission Planning Databook as a requirement to meet the Entry Descent Landing (EDL) target condition to a performance of better than 95 percent is the Near Simultaneous Tracking (NST). Additional data types discussed were the use of a three-way measurement and a difference range process. These data types would be used independently to assess the two-way coherent measurement data types (range and Doppler) baselined by the prime operations navigation team. During the presentations to the Mishap Investigation Board, it was stated that the Mars Polar Lander navigation team lead would be involved in the detailed analysis of the NST data. The application of a NST data type is relatively new to the Mars Polar Lander mission navigation procedure. These data types have not been previously used for Mars Climate Orbiter or Mars Polar Lander navigation. The results of the new data types in addition to range and Doppler only-solutions could potentially add to the uncertainty of the best estimate of the trajectory at the EDL conditions." [564]

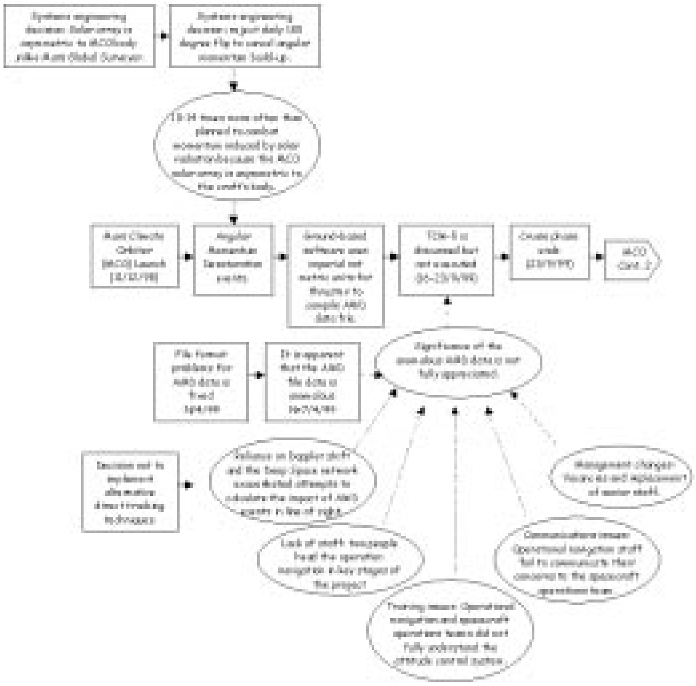
Mars Climate Orbiter (MCO) Launch
Angular Momentum Desaturation events
TCM-5 is discussed but not executed
Cruise phase ends
MCO Cont. 2
Significance of the anomalous AMD data is not fully appreciated
Decision not to implement alternative direct tracking techniques

Figure 10.12: Technological Barriers Fail to Protect the Climate Orbiter

Figure 10.12 introduces these technological issues into previous ECF diagrams. This diagram includes an event labelled Decision not to implement alternative tracking techniques and a condition

Reliance on Doppler shift measurements and the Deep Space network exacerbated attempts to directly observe the impact of AMD events. As can be seen, this reliance upon a particular tracking technology contributed to the failure of the people-based barriers mentioned in previous sections. This analysis raises a number of additional meta-level points that can be made about the use of barrier analysis to drive the development of ECF charts. It introduces a new event into the primary sequence. This denotes the decision not to initiate the TCM-5 maneuver. Although we have distinguished between the people, process and technology-based barriers, incidents often stem from complex interactions between these different protection mechanisms. A failure in one area of a system, as we have often seen, will compromise other forms of protection. The difficulties of making direct observations about the AMD events frustrated attempts to quantify any residual navigation error. The significance of any such error was not fully understood; key personnel were not familiar with the Climate Orbiter's operating characteristics.

Previous paragraphs have used a relatively high-level barrier analysis to refine and guide the development of more detailed ECF charts. For example, Table 10.1 is relatively abstract when compared with the more detailed events and conditions in Figure 10.12. It is, however, possible to construct barrier tables that capture more detailed observations about the problems that exacerbate mission failures. Table 10.2 builds upon the previous analysis to look at the more detailed reasons why the software bugs in the trajectory modelling were propagated beyond the development of the Climate Orbiter. These reasons focus on three potential barriers. The Software Interface Specification describe the units that were to be used within the project. In order to understand the failure of the Climate Orbiter, it is important to understand why this specification was not followed. The development and operations team also had detailed plans for the validation of system components. Again, it is important to understand why these plans failed to ensure the success of the mission. Finally, JPL supported a form of incident reporting system known as the Incident, Surprise, Anomaly scheme. This was deliberately intended to ensure that concerns, such as the anomalous data from the ground navigation software, was not ignored. If it had been reported to the system, there is a good chance that the concerns of the navigation team would have been addressed before TCM-5.

Hazard: Impact/Re-Entry	Target: Mars Climate Orbiter
Level 2 Technology: Incorrect Trajectory Modelling	
Barrier	Reason for failure?
Software Interface Specification	No software audit to ensure SIS conformance
	Poor navigation-spacecraft team communication.
	Inadequate training on importance of SIS
Software Testing and Validation	Unclear if independent tests conducted.
	Failure to recognise mission critical software.
	Poor understanding of interface issues
Incident Reporting Systems	Team member did not use ISA scheme.
	Leaders fail to encourage reporting.
	Domain experts not consulted.

Table 10.2: Level 2 Barrier Table for the Loss of the Climate Orbiter.

The Mars Surveyor Operators Project was guided by a Software Interface Specification (SIS) that both the format and units of the AMD file. This file was generated by SM_FORCES software running on ground-based computers. In order to satisfy the SIS requirements it was anticipated that this software would use metric units of Newtons per second to represent thruster performance data. As we have seen, however, the SM_FORCES software used English units of pounds per second. Subsequent processing of the AMD data by the navigation software algorithms therefore, underestimated the effect of AMD events on the spacecraft trajectory. The data was incorrect by a factor of 4.45; the ratio of force in pounds to Newtons. The SIS was intended to provide an important barrier against the type of software problems that led to the navigation software error.

The previous analysis does not, however, explain why the SIS failed to protect the system in the manner intended. Primary and secondary investigations identified inadequate training a key reason why development engineers failed to satisfy the interface requirements: “the small forces software development team needed additional training in the ground software development process and in the use and importance of following the Mission Operations SIS” [564].

Inadequate training about the importance of the SIS was compounded by a lack of training about appropriate testing techniques for the ‘small forces’ software. Not only did this increase the likelihood that the software would not comply with project interface requirements but it also reduced the likelihood that any anomalies would be identified. The investigators expressed a number of additional concerns about the testing procedures that were used during the development of the Climate Orbiter. It was unclear whether or not the ground software had been inspected by an independent validator. This lack of rigour can be explained by a possible perception that the small forces software was not ‘mission critical’. It can, therefore, be argued that the technological defences of an independent verification and validation program were breached by a managerial lack of oversight and the decision not to perform a system level hazard analysis.

The Mishap Board recommended that the Polar Lander teams should develop a verification matrix. One axis would denote all mission-critical project requirements. A second axis would denote the subsequent ‘mile-posts’ in mission development. A cell in the table would only be ticked if developers could present test results to demonstrate that the associated requirement had been met. The intention was that the verification matrix would explicitly record the test results for various requirements in Interface Control Documents, such as the SIS. It was also argued that the technical end-users of ground software applications should be required to sign-off these verification matrices.

Previous paragraphs have argued that limited training of key development staff led to an ignorance about the SIS and to inadequate testing of ground based software, including the small forces routines. Inadequate training also compromised a number of other barriers that might have protected the Climate Orbiter. In particular, the secondary investigation found members of the project team that did not understand the purpose or mechanisms of the Incident, Surprise, Anomaly (ISA) scheme. This finding is particularly important given the topic of this book. The ISA system was the primary means of providing information about adverse occurrences. Potential faults were logged with the system. Any subsequent remedial actions were then carefully monitored to ensure that the underlying issues were dealt with:

> “A critical deficiency in Mars Climate Orbiter project management was the lack of discipline in reporting problems and insufficient follow-up. The primary, structured problem-reporting procedure used by the Jet Propulsion Laboratory the Incident, Surprise, Anomaly process was not embraced by the whole team. Project leadership did not instill the necessary sense of authority and responsibility in workers that would have spurred them to broadcast problems they detected so those problems might be articulated, interpreted and elevated to the highest appropriate level, until resolved.” [570]

It is difficult to underestimate the importance of these points. If the navigation anomalies has been reported to the ISA system then there is a good chance that the navigation and spacecraft operations teams would have been requested to provide a coordinated response. This response might also have involved mission scientists who had the most knowledge of Mars, of the on-board instruments and of the mission science objectives. The investigators subsequently argued that their input could well have reversed the decision not to perform the TCM-5 maneuver.

Figure 10.13 presents an ECF chart that captures some of the more detailed events and conditions that helped to undermine the defences against software ‘bugs’ on the Climate Orbityer mission. As can be seen, the insights provided by the previous barrier analysis relate to two different stages in the mission. The top-left of the diagram represents the developers’ failure to use the SIS or then to discover that this interface had been violated. Events have been introduced to represent that the SM_Forces routines are written using imperial and not metric units for thruster performance and that Limited independent testing of the ground based SM_Forces routines took place. In contrast, the lower left-hand side of Figure 10.13 represents the failure of the operational staff to report the apparent navigation anomaly using the ISA scheme.

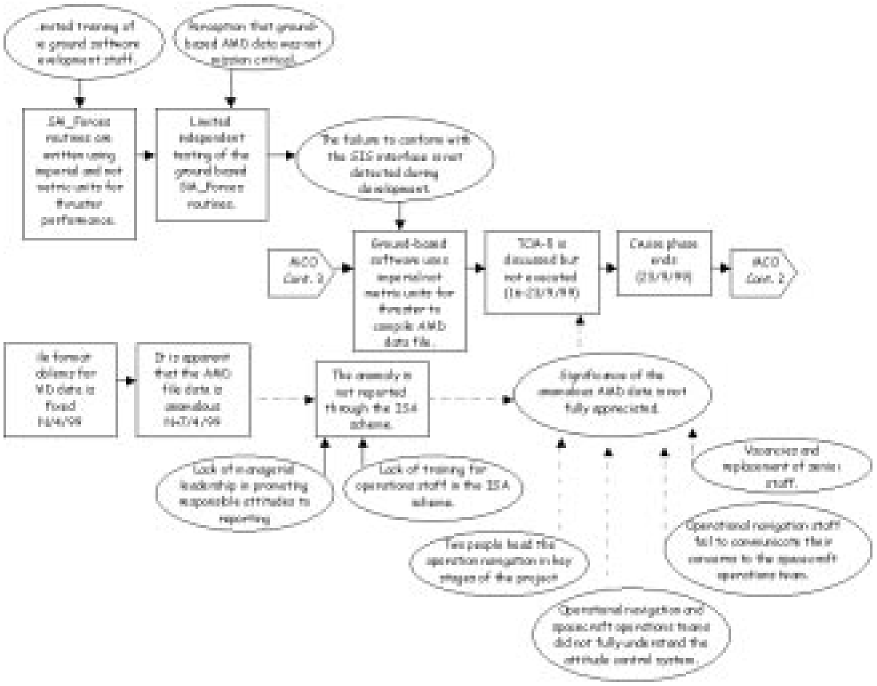
..imited training of ..e ground software ..evelopment staff.
Perception that ground-based AMD data was not mission critical.
SM_Forces routines are written using imperial and not metric units for thruster performance.
Limited independent testing of the ground based SM_Forces routines.
The failure to conform with the SIS interface is not detected during development.
MCO Cont. 3
Ground-based software uses imperial not metric units for thruster to compile AMD data file.
TCM-5 is discussed but not executed (16-23/9/99)
Cause phase ends (23/9/99)
MCO Cont. 2
..le format ..ciency for ..MD data is fixed 16/4/99
It is apparent that the AMD file data is anomalous 14-7/4/99
The anomaly is not reported through the ISA scheme.
Significance of the anomalous AMD data is not fully appreciated.
Lack of managerial leadership in promoting responsible attitudes to reporting
Lack of training for operations staff in the ISA scheme.
Vacancies and replacement of senior staff.
Two people lead the operation navigation in key stages of the project
Operational navigation staff fail to communicate their concerns to the spacecraft operations team.
Operational navigation and spacecraft operations teams did not fully understand the attitude control system.

Figure 10.13: Technological Barriers Fail to Protect the Climate Orbiter (2)

As can be seen, training failures are represented by conditions in both areas of this diagram. This observation has a more general significance beyond our analysis of the Climate Orbiter mission. Chapter 3 argued that training is often perceived to be a low cost work-around for a range of deeper design, development and management problems. It should not, therefore, be surprising if inadequate training is often identified in the role of a failed barrier or inadequate form of protection. It is regrettable that 'improved training' is often advocated as the remedy for this problem. More might be gained from a closer examination of why training failed to provide necessary protection in the first place.

10.2.3 Change Analysis

Previous section have shown how barrier analysis can direct the construction of ECF diagrams. Previous sections have not, however, shown that ECF diagrams can be used to distinguish between root causes and contributory factors. This is a deliberate decision. As we shall see, investigators must consider a range of information about the course of an incident before attempting such a causal analysis. The following paragraphs, therefore, present a further techniques that can be used to identify further information that can the be used to identify the root causes of an incident. Rather than repeat a barrier analysis for the Polar Lander incident, this section shows how change analysis can also be used as a precursor to this causal interpretation of an adverse occurrence.

The US Department of Energy [207], Occupational Safety and Health Administration (OSHA) [649] and NASA [571] all advocate change analysis as a key analytical tool for incident investigation. This technique was pioneered by Johnson in the year immediate after the Second World War. It was then developed for use by the US Airforce by Kepner and Tregoe in the Rand Corporation [248]. Change analysis can be used to determine whether or not abnormal working practices contributed to the causes of an adverse occurrence. The focus of this analytical technique is justified by the observation that deviations from normal operations are often cited as a cause in many accidents and incidents [207]. It is important to emphasise, however, that these changes are often made with the best intentions. For instance, new working practices may help to ensure that organisations satisfy regulatory requirements. Alternatively, new production processes can be introduced to improve organisational efficiency. Problems arise not from the intention behind such changes but from the difficult of predicting the impact that even small changes can have upon the operation of complex, technological systems. Even apparently beneficial changes can have unintended consequences that, in the medium or long term, can help to produce incidents and accidents.

In incident investigation, change analysis can be applied to identify the differences between what was expected to occur and what actually did occur during. OSHA's guidelines for incident and accident investigation include a brief tutorial on change analysis [649]. The following list enumerates the key stages in the OSHA approach. The US Department of Energy omit the final two stages and, instead, argue that investigators should feed the results of any change analysis into techniques that are intended to distinguish root causes from contributory factors [207]. They recommend that these findings should inform the development of the ECF charts, introduced in this chapter:

1. Define the problem.

2. Establish what should have happened?

3. Identify, locate and describe the change.

4. Specify what was and what was not affected.

5. Identify the distinctive features of the change.

6. List the possible causes.

7. Select the most likely causes.

Both the Department of Energy and OSHA provide relatively high-level guidelines for the application of change analysis. This is important because they provide investigators with an overview of the key stages that contribute to this technique. Unfortunately, these high-level summaries can also hide some of the underlying problems that complicate change analysis within many incident investigations. For instance, it is not always easy to determine what ought to happen during normal operation. The Polar Lander and Climate Orbiter missions had many unique characteristics that made them very different from similar projects. On the other hand, it is unclear whether or not it is possible to define what might be expected to happen during a normal NASA mission. The pressure to use leading-edge technology in pursuit of heterogeneous scientific objectives makes each mission very different from the last. Even in systems that have a greater 'routine', it can be difficult to identify operating norms. For example, the Department of Energy guidelines suggest that investigators use blueprints, equipment description documents, drawings and schematics, operating and maintenance procedures, job/hazard analyses, performance indicators etc to determine the nominal operating conditions before any incident [207]. However, subtle differences often distinguish the ways in which different plants operate the same process. Even within a plant, there will be differences in the performance of different shifts and of individuals within those shifts. Similarly, the notion of an accident-free or ideal situation can be difficult to sustain in many industries. For instance, some oil installations operate running maintenance programs. Temporary fixes are used to resolve non-critical failures. This enables operations to continue until a scheduled maintenance period. This interval is used to conduct longer-term repairs. Such maintenance schemes raise a number of questions about what is, and what is not, a nominal state. For instance, operators view the system as operating normally even though it requires longer-term maintenance. This may seem to be an isolated example. This argument can, however, be applied to a more general class of systems. Most applications continue to operate in spite of documented failures in non-critical components. Some authors have gone further and argue that complex, safety-critical systems are unlikely to be error-free [675]. They always involve adaptations and work-arounds because it is impossible for designers and operators to predict the impact that the environment will have upon their systems.

Further problems stem from the effects of compound changes. For example, operating practices and procedures evolve slowly over time so that official documents may reflect a situation that held several years previously. Under such circumstances, previous distinctions between normal and abnormal practices can become extremely blurred. Other problems arise when changes that occurred several years before are compounded by more recent changes. The change analysis guidelines suggest that investigators should address such situations by developing several baseline or nominal situations. The events during an incident should be contrasted with normal working practices immediately prior to any failure and also with normal working practices in the years before to any previous change:

> "...decreases in funding levels for safety training and equipment may incrementally erode safety. Compare the accident scenario to more than one baseline situation, for example one year ago and five years ago, then comparing the one and five year baselines with each other can help identify the compounding effects of change."[207]

Chapters 6 and 7 have already described the difficulties that can arise when investigators must piece together the events that contribute to a particular incident. Automatic logging systems can be unreliable and seldom capture all critical aspects of an adverse occurrence. It can also be difficult to interpret the information that they do capture. Individuals may be unable to recall what happened in the aftermath of an adverse occurrence. In the aftermath of an incident, there is also a temptation for operators to describe violations as abnormal occurrences even though they may have formed part of everyday working practices. Organisation, managerial and social pressures influence their participation in a primary and secondary investigation. Inconsistencies, omissions and ambiguity are a continual problem when investigators must form coherent accounts from eye-witness statements. All of these factors combine to frustrate attempts to determine ways in which an incident differed from 'normal' practice. Change analysis must also consider a number of further issues. It is usually insufficient simply to contrast normal behaviour with the abnormal events that occur during an incident. One an incident has occurred, it is also important for investigators to determine the success or failure of any remedial or mitigating actions. Given that an incident occurred, it is important to determine whether or not the response followed pre-determined procedures.

These caveat are important because they identify some of the practical difficulties that emerge during the application of change analysis. It is also important to notice, however, that they do not simply affect this analytical technique. The problems of eliciting evidence and reconstructing an incident are common to all incident investigation. Change analysis is unusual because it forces investigators to explicitly address these issues during their analysis. Other techniques, including barrier analysis, make no distinction between the normal and abnormal events that contribute to an incident.

Meta-Level Change Analysis

Reason [701] argues that incidents and accidents often stem from underlying changes in the structure of complex organisations. Change analysis can, therefore, begin in a top-down fashion by considering the organisational context in which the Polar Lander mission took place. In particular, it is important to consider the consequences of the "Faster, Better, Cheaper" strategy that was introduced by the NASA Administrator, Daniel Goldin. He assumed command at a time of shrinking financial resources caused by the recession of the early 1990's. The US government had responded to global economic problems with a program of deficit reduction that affected many including education, healthcare and housing. Golding was faced by a situation in which NASA was likely to receive insufficient funds to cover all of its future programme commitments. He, therefore, conducted a thorough review of both existing and future projects using 'red' and 'blue' teams. These groups were to analyse both the programmes themselves and their organisational context. Blue teams examined their own programs for creative ways to reduce cost without compromising safety or science. Red teams were composed of external assessors who were intended to bring in new ideas and to ensure that those ideas were realised. This review began in May 1992 and had an almost immediate impact. By December 1992, it was claimed to have delivered a seventeen percent reduction in costs [576].

The cost improvements and efficiencies that were achieved under the new "Faster, Better, Cheaper" initiative had a profound impact on the relationship between NASA and its contractors. As we shall see, changes in this relationship were at the heart of the problems experiences during the Climate Orbiter and the Polar Lander missions. In particular, an Independent Cost Assessment Group was set up to ensure that cost estimates were as accurate as possible. This followed a General Accounting Office report into a sample of 29 NASA programs that identified an average cost growth of 75 percent. Goldin argued that "We can not tolerate contracts so fluid, that the product we bargained for in no way resembles what we end up with... We are partners with industry, but we will hold you [contractors] accountable for what you sign up to deliver and ourselves accountable for establishing firm requirements" [577].

It is difficult to find a precise definition of what the "Faster, Better, Cheaper" initiative was supposed to imply at a project level. The Mars Program Independent Assessment Team was formed after the loss of the Polar Lander [569], it identified the following components of this initiative:

- *Create smaller spacecraft for more frequent missions.* The creation of smaller, more frequent missions was intended to increase the opportunities for scientists, and the public, to participate in NASA's work. This approach was also perceived to have the additional benefit of distributing risk across the increased number of projects. The "Faster, Better, Cheaper" strategy distributes the risk of achieving science objectives among more missions thus minimising the impact of a single mission failure;

- *Reduce the cycle time throughout a project.* Increased mission frequency was intended to help introduce scientific and engineering innovations. This would be achieved by reducing project lead time. Such reductions were not be made by the arbitrary curtailment of development or implementation time. They were to be achieved by the elimination of inefficient or redundant processes and, especially, through the use of improved management techniques and engineering tools In the Polar Lander and Climate Orbiter missions, this involved greater responsibilities for line management within individual project contractors;

- *Use new technology.* The "Faster, Better, Cheaper" strategy relied upon the integration of new technology into many different aspects of each mission. New technology was intended both to increase the scientific return of each mission, to reduce spacecraft size and to limit overall mission cost. It was, however, recognised that new technologies must "be adequately mature" before being incorporated in a flight program [569]. This use of innovative technology was also intended to increase public interest in NASA programs;

- *Accept prudent risk if they are warranted by the potential rewards.* It was recognised from its inception that the "Faster, Better, Cheaper" implied taking risks; "in all cases, risks should be evaluated and weighed against the expected return and acknowledged at all levels" [569]. Rather than using flight-proven techniques, programs were encouraged to incorporate new technologies if they showed promise of significantly increasing mission capabilities or improving efficiency. The use of the term 'prudent' in many of the "Faster, Better, Cheaper" documents was intended to ensure that these technologies underwent a rigorous testing and validation prior to their use in flights. This was encapsulated in the maxim 'Test-As-You-Fly/Fly-As-You-Test'; validation should provide a close approximation of the eventual mission characteristics.

- *Use proven engineering and management practices to maximise the likelihood of mission success.* The technological risks associated with this new strategy were to be addressed using proven engineering and management techniques. These techniques were to include hazard analysis, using Fault Tree Analysis or Failure Effects and Criticality Analysis. There was an explicit concern to prevent any 'single human mistake causing mission failure' [569]. These established techniques were also to establish a chain of responsibilities and reporting within each project. Projects were to be reviewed by independent experts from outside the projects or implementing institutions. These individuals were to provide an overall project assessment and to review any associated risks.

This description of the "Faster, Better, Cheaper" strategy acts as a statement of what was intended by Administrator Goldin's initiatives. It, therefore, provides an ideal or standard against which to compare the particular characteristics of the Polar Lander project. This is important given the specialised nature of such missions, change analysis has most often been applied to process industries that follow more regular patterns of production. Table 10.3, therefore, uses this approach to assess the differences between the intended objectives of the "Faster, Better, Cheaper" strategy and what went on during the Mars Surveyor'98 projects. In particular, it summarises the investigators argument that the Polar Lander team were forced to:

> "Reduce the cost of implementing flight projects in response to severe and unprecedented technical and fiscal constraints... One lesson that should not be learned is to reject out of hand all the management and implementation approaches used by these projects to operate within constraints that, in hindsight, were not realistic." [579]

It is important to emphasise that Table 10.3 does *not* compare the Polar Lander mission with missions that took place before the Goldin initiative. Such a comparison would be academically interesting but might also ignore the changing financial circumstances that have fundamentally changed the way that NASA operates in recent years.

Prior/Ideal condition	Present Condition	Effects of change
Faster, better, cheaper strategy required sufficient investment to validate high-risk technologies before launch	Mars Surveyor'98 faces pressures to push boundaries of technology and cost	Greater development effort
		Use off-the-shelf hardware and inherited designs as much as possible.
		Use analysis and modeling as cheaper alternatives to system test and validation.
		Limit changes to those required to correct known problems; resist changes that do not manifestly contribute to mission success.

Table 10.3: High-Level Change Table for the MPL Mission.

The first entry in Table 10.3, therefore, summarises the intended effects of the "Faster, Better, Cheaper" strategy on the Polar Lander mission. In contrast, NASA's investigators found evidence to suggest that the Mars Surveyor projects pushed the limits of what was possible both technologically and within available budgets. The pressure to push the technological boundaries are illustrates by the Deep Space 2 probes. These were designed to test ten high-risk, high-payoff technologies as part of NASA's New Millennium Program. They were to demonstrate that miniaturised components could be delivered to the surface of another planet and could be used to conduct science experiments. The risks associated with this new technology were assessed and approved by JPL and NASA management [579]. The risk-assessment was, however, performed on the assumption that there would be a ground-based system-level, high-impact test. This test was not conducted because of budgetary constraints. Although this is a specific example, it supports the higher level observation in Table 10.3 that the Surveyor projects pushed the boundaries both of technology and cost. A further illustration can be provided by a comparison between the Mars Surveyor'98 missions and the previous Pathfinder project. Pathfinder demonstrated the successful application of a comparable range of technological innovation under the "Faster, Better, Cheaper" strategy. NASA have, however, estimated that the

Mars Surveyor missions were underfunded by up to 30% in comparison with the Pathfinder [569]. This estimate is supported by the funding summary in Table 10.4.

	Pathfinder	Mars Surveyor'98 (MCO and MPL)
Project Management	11	5
Mission Engineering and Operations Development	10	6
Flight System	134	133
Science and Instrument Development	14	37
Rover	25	0
Other	2	7
Total	196	188

Table 10.4: Comparison of the Development Costs for the Pathfinder and Mars Surveyor'98 (in $ Millions at 1999 prices).

Table 10.3 summarises the impact that budgetary pressures had upon the technological development of the Polar Lander. Developers made a number of decisions that were based on budgetary considerations but which ultimately had a critical effect upon systems engineering. These included decisions to use off-the-shelf components and inherited designs as much as possible. Analysis and modeling were also to be used as lower-cost alternatives to system test and validation. Changes were to be limited to those required to correct known problems. There was pressure to resist changes that did not directly contribute to mission success. The following sections look beyond these high level effects. Change analysis is used to analyse the detailed engineering and managerial impact of the Polar Lander's "Faster, Better, Cheaper" objectives. The results of this analysis are then used to inform the ECF charts that were presented in Figures 10.6 and 10.7.

In passing, it is worth noting that Table 10.3 illustrates some of the limitations of change analysis at this relatively high level of abstraction. It does not explain the reasons why the Surveyor'98 project adopted this extreme version of Goldin's policy. Subsequent investigations argued that this was due to ineffective communication between JPL management and NASA Headquarters. NASA Headquarters thought it was articulating program objectives, mission requirements, and constraints. JPL management interpreted these statements as non-negotiable program mandates that specified particular launch vehicles, costs, schedules and performance requirements [569].

Figure 10.14 illustrates the way in which the findings from an initial change analysis can be integrated into a high level ECF chart. This is a relatively straightforward process because the *present condition* in a Change Analysis, such as Table 10.3, can be directly introduced as a condition within an ECF chart. In Figure 10.14 this is denoted by the note that is labelled Mars Surveyor'98 faces pressures to push boundaries of cost and technology. The change analysis does not, however, identify which events this present condition will effect within an ECF chart. The node labelled Launch approved has, therefore, been introduced into Figure 10.14. Later sections will refine this high-level event to look at a number of specific events that were affected by the Faster, Better, Cheaper strategy. The change analysis illustrated in Table 10.3 also documented a number of effects that stem from the higher-level pressures to innovate and cut costs. For example, previous paragraphs have mentioned the policy to exploit off-the-shelf hardware and inherited designs as much as possible. These effects cannot be introduced directly into ECF charts. As we shall see, they occasionally refer to particular events. In this instance, they denote more specific conditions that influence the events leading to the loss of the Polar Lander. This illustrates the important point that analysts must still interpret and filter the information that is obtained using techniques such as change and barrier analysis. These is not automatic translation between the information that is derived from these approaches and their graphical representation in an ECF chart.

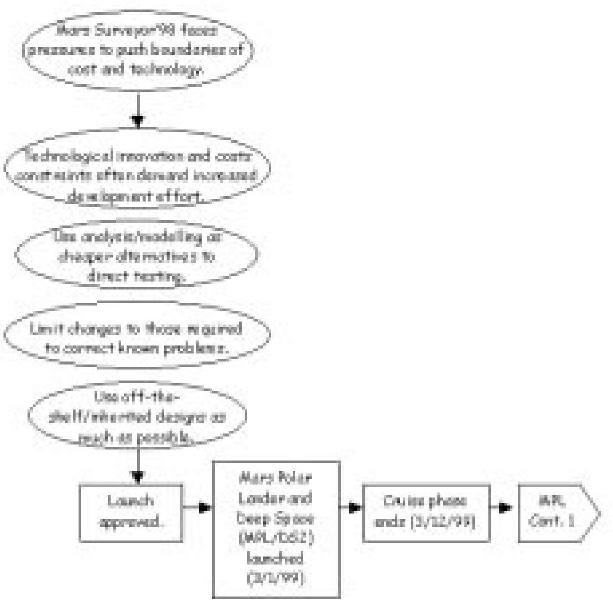
Mars Surveyor'98 faced pressures to push boundaries of cost and technology.
Technological innovation and costs constraints often demand increased development effort.
Use analysis/modelling as cheaper alternatives to direct testing.
Limit changes to those required to correct known problems.
Use off-the-shelf/inherited designs as much as possible.
Launch approved.
Mars Polar Lander and Deep Space (MPL/DS2) launched (3/1/99)
Cruise phase ends (3/12/99)
MPL Cont. 1

Figure 10.14: Integrating Change Analysis into an ECF Chart

People: Changes in Staffing Policy

One aspect of the "Faster, Better, Cheaper" strategy was that NASA was to profit by a greater involvement with commercial organisations. The intention was to retain a civil service and JPL core competency for in-house science, research and engineering. Aerospace operations, including the operation of the Space Shuttle and the Surveyor program, were to be performed by NASA contractors. There was also a plan to transfer program management responsibility to the field Centers from NASA Headquarters. The 1996 budgetary statement also included a commitment to performance-based contracting:

> "$100 million savings are presently projected as a result of implementing performance-based contracts for aeronautical research and facility maintenance and operations. The savings come from reducing contractor staffing levels by asking the contractor to use their ingenuity in carrying out the required work. NASA will specify what we want and when it is needed vs. specifically directing the contractor not only what and when, but also how to do the job. This will involve conversion of many current NASA cost-reimbursement/level-of-effort, specification-laden contracts." [561]

As we shall see, this contractor 'ingenuity' helped to erode a number of important safety mechanisms in order to meet the relevant budgetary constraints. Contractor staff habitually worked excessive amounts of overtime. There was often only a single expert available within key mission areas.

Table 10.5 summarises the differences between the planned use of contract management and the experience of the Polar Lander mission. The intention was to reduce costs by relying on the contractor's existing management structure to run the day to day operation of the project. The ten or so JPL staff who were involved in the project were primarily intended to provide higher-level oversight. This was a departure from previous JPL projects and the result was minimal involvement by JPL technical experts.

It is worth reiterating that the project team was expected to deliver a lander onto the surface of Mars for approximately one-half of the cost of the Pathfinder mission. Under such constraints, it was difficult for the contractor's staff to meet their commitments within the available resources. LMA used excessive overtime in order to complete the work on schedule. Many development staff worked for sixty hours per week [579]. Some worked more than eighty hours per week for extended

Prior/Ideal condition	Present Condition	Effects of change
Greater JPL line-management involvement in the project.	LMA staff found it hard to fulfill mission requirements with available resources.	LMA used excessive overtime to complete work on schedule.
		Many key technical areas were staffed by a single individual.
		Lack of peer interaction.
		Breakdown in intergroup communications.
		Insufficient time to reflect on unintended consequences of day-to-day decisions.
		Less checks and balances normally found in JPL projects.

Table 10.5: Change Summary Table of MPL Staffing Issues.

periods of time. Budgetary constraints created further technical problems because key areas were only staffed by a single individual. This removed important protection mechanisms because it became difficult to arrange the continual peer review and exchange of ideas that had characterised previous projects. The workload may also have jeopardised communications between technical disciplines. There was insufficient time and workforce available to provide the checks and balances that characterised previous JPL missions.

Figure 10.15 provides a further illustration of the way in which change analysis can be used to inform the construction of an ECF chart. As can be seen, the additional analysis of staffing issues has helped to identify a number of conditions that affected both the development and the subsequent validation of the lander's design. As a result, the higher-level conditions that were identified in Figure 10.14, such as use analysis/modelling as cheaper alternatives to direct testing, have been reorganised into the three strands shown in Figure 10.15. These strands distinguish between conditions that relate narrowly to staff limitations, such as the use of single individuals to cover key technical areas, from wider issues relating to the technological demands and validation of projects under the faster, better, cheaper strategy. This illustrates another important point about the process of integrating the findings of barrier and change analysis into ECF charts. The introduction of new information can force revisions to previous versions of the diagram. These revisions may result in conditions or events being removed, merged, edited or moved.

Figure 10.15 introduces a further extension to the ECF notation. A horizontal parenthesis is used to indicate that conditions from a high-level change analysis and an analysis of staffing issues influence both the development and the launch approval process. Subsequent analysis might avoid this additional syntax by omitting one of the first two events in this diagram. This has not been done because some conditions, such as the lack of peer interaction, may not only have affected the decision to launch but also the development process that led to that event. Alternatively this additional syntax could be omitted if conditions were assigned to either the development or the launch approval events. For example, the use of analysis and modelling rather than direct testing might be associated with the decision to launch rather than the completion of the development phase. Such distinctions seem to be arbitrary and have, therefore, been avoided.

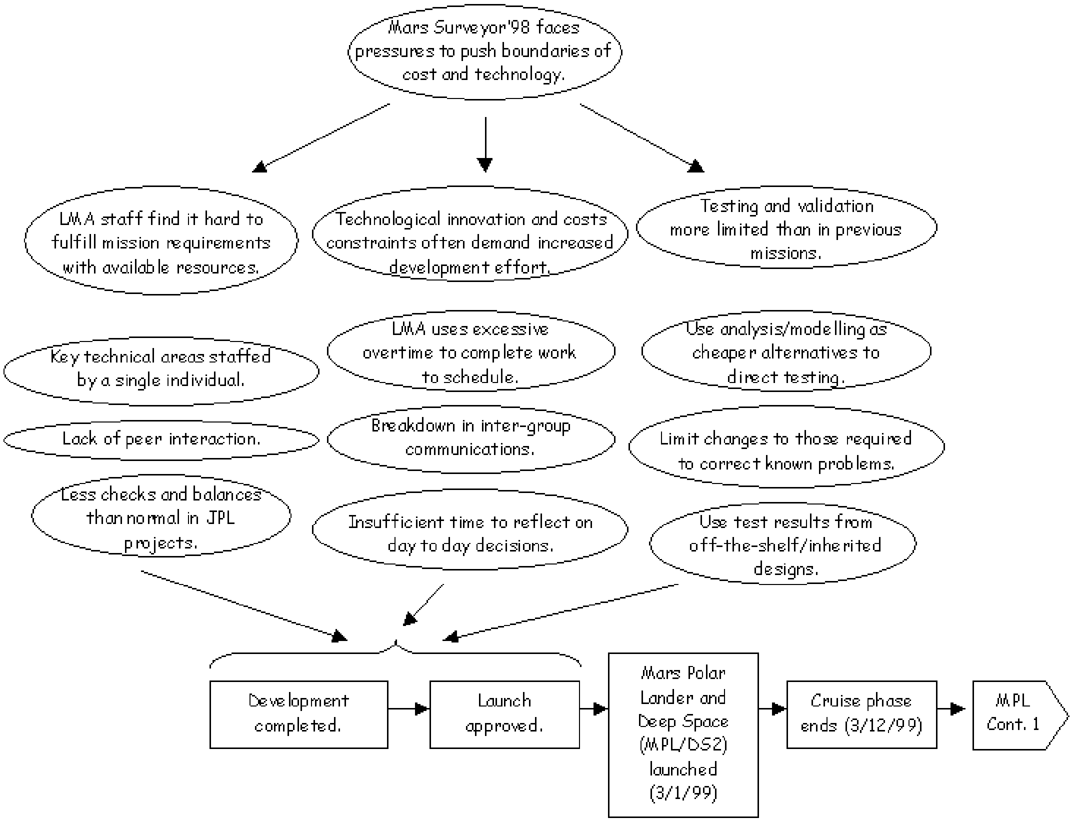
Mars Surveyor'98 faces pressures to push boundaries of cost and technology.
LMA staff find it hard to fulfill mission requirements with available resources.
Technological innovation and costs constraints often demand increased development effort.
Testing and validation more limited than in previous missions.
Key technical areas staffed by a single individual.
LMA uses excessive overtime to complete work to schedule.
Use analysis/modelling as cheaper alternatives to direct testing.
Lack of peer interaction.
Breakdown in inter-group communications.
Limit changes to those required to correct known problems.
Less checks and balances than normal in JPL projects.
Insufficient time to reflect on day to day decisions.
Use test results from off-the-shelf/inherited designs.
Development completed.
Launch approved.
Mars Polar Lander and Deep Space (MPL/DS2) launched (3/1/99)
Cruise phase ends (3/12/99)
MPL Cont. 1

Figure 10.15: Representing Staffing Limitations within an ECF Chart

Technology: Changes in Innovation and Risk Management

A number of consequences stemmed from these changes in the staffing of the Polar Lander project. In particular, the communications problems that were noted by the investigators may have compromised necessary hazard analysis. In order to assess the impact of this, it is again important to establish NASA policy for an 'ideal' approach to risk management:

> "To reduce risk, we need to manage our projects systematically, especially if we expect to be successful with faster, better, cheaper projects. The Risk Management process efficiently identifies, analyses, plans, tracks, controls, communicates, and documents risk to increase the likelihood of achieving program/project goals. Every project should have a prioritized list of its risks at any point in the life cycle, along with the programmatic impacts. The list should indicate which risks have the highest probability, which have the highest consequences, and which need to be worked now. It means that all members of the project team should have access to the risk list so that everyone knows what the risks are. It means that the project team members are responsible for the risks. The team should work to reduce or eliminate the risks that exist and develop contingency plans, so that we are prepared should a risk become a real problem... From the beginning of a project, the Project Manager and team should have an idea of what the 'risk signature' of the project will be. The risk signature will identify expected risks over the course of the project and when the project risks are expected to increase and decrease. During the project, risks should be tracked to determine if mitigation efforts are working. " [573]

This policy is promoted through a range of publications and courses that are supported by NASA's Office of Safety and Mission Assurance. Change analysis again provides a means of contrasting these 'ideals' with the experience of the Polar lander project. Table 10.6 provides a high level view of the differences that emerge.

Prior/Ideal condition	Present Condition	Effects of change
Adequate risk assessment at system level	No system-level Fault Tree analysis was formally conducted or documented	Bottom-up Failure Modes, Effects and Criticality Analysis hides higher-level interaction/systemic issues
		No risk analysis of propulsion, thermal and control interaction.
Adequate risk assessment at subsystem level	Fault-tree analysis treated inconsistently for different subsystems	Bug in timer for uplink loss found in Fault Tree after loss of flight.
		Premature trigger of touchdown sensor found in Fault Tree before Entry, Descent and Landing but not guarded against.
Project management maintains explicit risk-signature for the project	No risk assessment for going beyond Preliminary Design Review with 15% mass margin.	Management focus on mass reduction not risk reduction activities.

Table 10.6: Change Summary Table of MPL Risk Management.

This table suggests that risk analysis should have been conducted in a systematic manner across the various subsystems but also at a project level. There was no explicit attempt to model the way in which system-level, mission, risks changed over time. NASA refers to this model as the risk signature of a project [579]. It is important because it provides managers with a means of tracking how particular development decisions can affect the risk-margins that are eroded by particular development decisions. For instance, the preliminary design review decided to proceed with only a 15% margin between the predicted mass of the Polar Lander and the capabilities of the chosen launch vehicle. This mass assessment also failed to account for a number of outstanding mass commitments. Previous projects might have anticipated a mass margin of at least 25%. This events illustrate how key decisions were informed by cursory risk assessments. The decision to proceed with a 15% mass margin also had a significant impact upon subsequent risk management. Project resources were diverted into mass reduction rather than risk reduction activities [579].

Failure Modes, Effects and Criticality Analysis (FMECA) was used to support many areas of systems engineering. This technique is, however, driven by a bottom-up analysis of failure modes. It cannot easily be used to analyse the interactions between complex sub-systems. System level properties are often lost when FMECA is used to analyse the failure modes of complex systems. Top-down risk analysis techniques can be used to overcome these limitations. A Fault Tree analysis was, therefore, conducted for specific mechanisms and deployment systems. This analysis was only conducted for those systems that were perceived to be particularly vulnerable, for instance, because they lacked any form of redundancy. As mentioned, there was no evidence of any system level fault tree analysis. In particular, there was an 'incomplete' analysis of the hazards that might emerge from the interaction between propulsion, thermal and control systems [579].

The problems of risk management not only affected the risk signature of the project and the hazards associated with subsystem interaction, further problems also affected individual subsystems. For example, there was a problem in the software that was designed to automatically re-establish communications links if the up-link was lost during the Entry, Descent and Landing phase. This bug was not detected before launch or during the cruise phase of the flight. A Fault Tree analysis identified this as a possible failure mode after the Polar Lander had been lost. This led to a more detailed examination of the code. External reviers were then used to validate the hypothesised failure. Even when risk management techniques did succeed in identifying a potential failure mode, sufficient actions were not always taken to ensure that the hazard could not arise. The Mission Safety and Success Team performed a fault-tree analysis of the Entry, Descent and Landing stage. The team then conducted an analysis to determine whether or not the design afforded sufficient protection against the identified hazard. They identified a potential failure if the Hall effect sensors received premature touchdown signals. This scenario is represented in Figure 10.7. They were, however, satisfied by the software design and testing that was provided by the contractors.

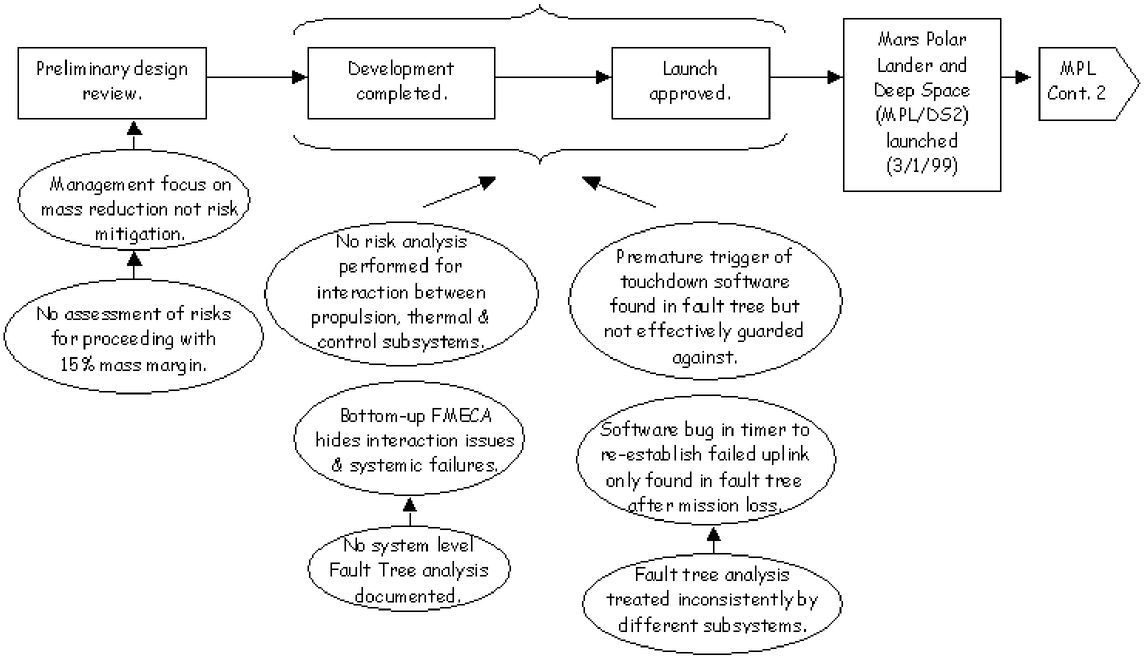
Preliminary design review.
Development completed.
Launch approved.
Mars Polar Lander and Deep Space (MPL/DS2) launched (3/1/99)
MPL Cont. 2
Management focus on mass reduction not risk mitigation.
No assessment of risks for proceeding with 15% mass margin.
No risk analysis performed for interaction between propulsion, thermal & control subsystems.
Bottom-up FMECA hides interaction issues & systemic failures.
No system level Fault Tree analysis documented.
Premature trigger of touchdown software found in fault tree but not effectively guarded against.
Software bug in timer to re-establish failed uplink only found in fault tree after mission loss.
Fault tree analysis treated inconsistently by different subsystems.

Figure 10.16: Representing Risk Management Issues within an ECF Chart

Figure 10.16 incorporates the insights from Table 10.6 into an ECF chart. The change analysis helps to identify some of the conditions that influenced events leading up to the loss of the Polar Lander. As before, some of these conditions affected many different aspects of the development process. These include the lack of any system level fault tree and the inconsistent way in which hazard analysis was performed within individual subsystems. Figure 10.16 also illustrates the way in which change analysis can be used at a more detailed level to assess the impact that departures from 'expected practice' had upon particular events. In particular, the lack of any assessment of the risks associated with proceeding on a mass margin of only 15% had a knock-on effect when management spent increasing amounts of time on mass reduction rather than risk mitigation. These two conditions are associated with the Preliminary Design Review. This event marks a critical stage when the projects mass margins are first established.

It is important to note that Figure 10.16 illustrates some of the limitations of the ECF notation. For example, the lack of any risk assessment for the 15% mass margins is associated with the

Preliminary Design Review. This condition had knock-on effects that influence many subsequent events. In particular, the managerial focus on mass reduction is shown in Figure 10.16 as affecting the Preliminary Design Review. It also clearly affected subsequent risk assessments. Unfortunately, this is difficult to denote within the existing ECF syntax. Such limitations have inspired researchers to investigate a host of more 'advanced' techniques. Some of these have been introduced in Chapter 9. It is, however, important to note the complexity of the situation that is being analysed. A condition, the lack of any risk analysis for the 15% margin, influenced an event, the Preliminary Design Review. The consequences of this event, and in particular the decision to proceed with a 15% margin, imposed conditions upon the rest of the development process, managers had to focus on mass reduction. Such situations could be denoted within the existing ECF syntax. Edges might be drawn between conditions and events that occur later in an incident sequence. This would, however, result in a proliferation of interconnections between conditions and events. Alternatively, a cross-referencing scheme might be introduced so that conditions could be repeated at different points within an ECF chart. It is worth emphasising that most analytical techniques suffer from similar problems. The process of scaling-up from small scale studies often leads to a point at which the notation fails to capture important properties of an incident. These problems can usually be addressed through accretions to the syntax and semantics of the notation. Unfortunately, this leads to problems in training others to use the new hybrid technique. This is a serious problem. Such notation extensions can only be justified if they provide benefits to 'real-world' incident investigators. Many notations have been developed and extended without any practical validation.

Previous sections have focussed on high-level changes in the way in which the Polar Lander mission was managed. In contrast, Table 10.7 assesses the impact of particular technological decisions. It is important to emphasise, however, that many of these decisions were motivated by higher-level management objectives. It is also important to emphasise that these objectives were extremely complex and, potentially, contradictory. On the one hand, budgetary constraints made it essential for NASA to justify it's expenditure on technological innovation. On the other hand, many previous missions exhibited an understandable conservatism based on the feeling that mission success could be assured through the use of proven technology. This conflicts can be clearly seen in the Federal review of NASA laboratories. This formed part of President Clinton's wider initiative that also examined the Department of Defence and Energy's facilities. The resulting report argued that NASA's relatively large scientific research budget produced "limited opportunities for developing technologies" to address the faster, better, cheaper strategy [572]. They also acknowledged, however, that the gap between technology development and technology utilization was the most significant problem faced by NASA's Space Technology Enterprise. The review also reported the strong tendency within NASA to incorporate only "flight-proven technology" into space-flight missions.

These diverse factors created unusual effects on the Polar Lander project. On the one hand, the Deep Space 2 project shows a strong desire to assess the capabilities of a range of technological innovation. On the other hand, the Lander itself was developed with the explicit intention of borrowing as much as possible from previously successful mission. The Polar Lander was equipped with a disk-gap-band parachute that was identical to the one used on the Pathfinder mission, except that the Pathfinder logo had been removed. It also used an Eagle-Picher type of battery from the same batch as the one used on Pathfinder. This overall policy was, however, compromised when developers identified potential opportunities to reduce the project budget. For example, the lander exploited off-the-shelf engines that forced revisions to the initial configuration. Such technical innovations met the objectives espoused by the proponents of faster, better, cheaper. They also increased the level of uncertainty associated with the Lander's eventual performance.

As mentioned, Table 10.7 summarises the consequences of pressures to exploit technological innovation as a means of supporting the faster, better, cheaper strategy. This assessment is supported by the NASA investigators. The investigators found that the decision not to have EDL telemetry was defensible in terms of the project budget. It was, however, indefensible in terms of the overall program because it placed severe constraints on the amount of information that could be gleaned from any potential failure. Finally, communications were compromised by the decision to base the Lander's X-band down-link on a medium gain antenna that had to be accurately pointed at the earth. There was no X-band down-link through the more 'forgiving' omni-antenna. This "reduced

Prior/Ideal condition	Present Condition	Effects of change
Throttle valve for descent engines.	Pulse-mode control.	More difficult terminal descent guidance algorithm.
Lander design based on 2 canted engines in 3 locations.	4 smaller off the shelf engines in 3 locations.	Additional design and validation complexity.
Entry, descent and landing telemetry is available	Entry, descent and landing telemetry was not available	Problems in determining causes of mishap to inform future of program.
Downlink possible through omni-antenna	X-band down-link dependent upon MGA being pointed accurately at Earth.	Reduced chance of obtaining engineering data after anomalous landing.

Table 10.7: Change Summary Table of MPL Technological Issues.

the ability to get health and safety engineering data in an anomalous landed configuration. [579]". The decision to use pulse-mode control for the descent engines avoided the cost and risk of qualifying a throttle valve. This, however, increased the complexity of the descent guidance algorithm and introduced further risks into the propulsion, mechanical, and control subsystems. The lander configuration required at least two canted engines in each of three locations for stability and control. The project elected to use four smaller off-the-shelf engines at each location.

Figure 10.17 again shows how the findings of a change analysis can be integrated into an ECF chart. In particular, this diagram focuses on the communications issues that restricted communication both during and immediately after the Entry, Descent and Landing phase of the mission. Table 10.7 captured the observation that, in retrospect, it would have been better to have provided telemetry data during Entry, Descent and Landing. The decision not to provide this facility was justified by the argument that "no resources would be expended on efforts that did not directly contribute to landing safely on the surface of Mars" [579]. As can be seen, Figure 10.17 represents this analysis as two conditions labelled Entry, descent and landing telemetry is not available and Problems in determining cause of mishap make it hard to identify lessons for future systems. These conditions are, in turn, linked to previous ECF charts by introducing an event that represents the establishment of the mishap board. Their work was complicated by the lack of telemetry data.

Figure 10.17 also includes conditions that represent the potential effects of a communication failure. This is done by the conditions that are labelled X-band down-link is dependent upon medium gain antenna being accurately pointed at Earth and Reduced chance of obtaining engineering data after anomalous landing. This raises a further problem in the application of ECF charts as a means of modelling complex incidents and accidents. Previous sections have mentioned that the lack of any telemetry data makes it difficult for investigators to be certain about the exact causes of the failure. In consequence, Figure 10.17 represents a scenario in which the Lander is lost through the software bug in the handling of spurious signals from the Hall effect sensors and the Deep Space 2 probes are lost from electrical failures at impact. If, however, the software bug did lead to the loss of the lander then the decision to rely on the Medium Gain Antenna for the X-band up-link becomes of secondary importance to this incident. The chances of the Lander surviving the resultant impact with the planet surface are so remote that it this decision would have had little effect on the incident. Figure 10.17, therefore, introduces a double-headed line to illustrate that the X-band link may be significant for other failure scenarios or for future missions but that it is of limited relevance to this incident.

Table 10.7 also summarises the inspectors argument that the limited budget created a number of problems in assessing the cost-risk tradeoff for particular technological decisions. The difficulty

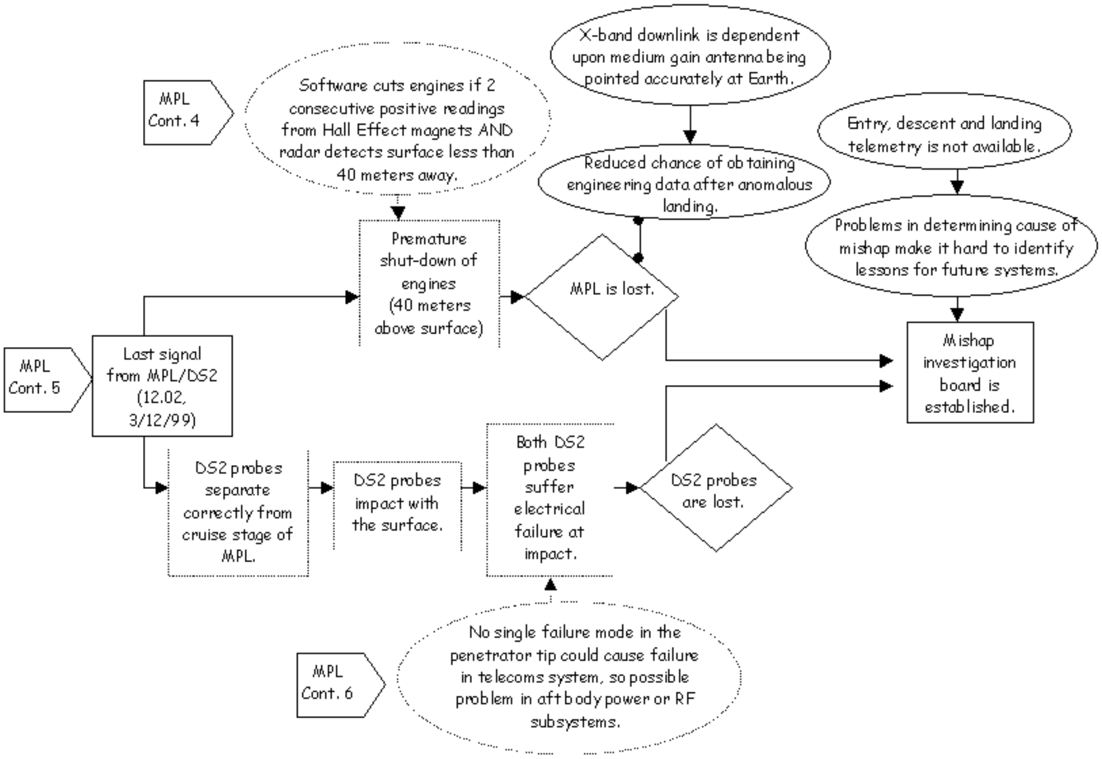
X-band downlink is dependent upon medium gain antenna being pointed accurately at Earth.
MPL Cont. 4
Software cuts engines if 2 consecutive positive readings from Hall Effect magnets AND radar detects surface less than 40 meters away.
Entry, descent and landing telemetry is not available.
Reduced chance of obtaining engineering data after anomalous landing.
Problems in determining cause of mishap make it hard to identify lessons for future systems.
Premature shut-down of engines (40 meters above surface)
MPL is lost.
Mishap investigation board is established.
MPL Cont. 5
Last signal from MPL/DS2 (12.02, 3/12/99)
DS2 probes separate correctly from cruise stage of MPL.
DS2 probes impact with the surface.
Both DS2 probes suffer electrical failure at impact.
DS2 probes are lost.
MPL Cont. 6
No single failure mode in the penetrator tip could cause failure in telecoms system, so possible problem in aft body power or RF subsystems.

Figure 10.17: Representing Technological Issues within an ECF chart (1)

of making such an assessment led to unanticipated design complexity. The decision to use pulse-mode control for the descent engines avoided the cost and risk of qualifying a throttle valve. This, however, increased the complexity of the descent guidance algorithm and introduced further risks into the propulsion, mechanical, and control subsystems. The lander configuration required at least two canted engines in each of three locations for stability and control. The project elected to use four smaller off-the-shelf engines at each location. Figure 10.18 represent two events in the development of the Lander: Decision to use pulse mode control and Decision to use off-the-shelf engines in 4x3 configuration. These events provide a specific example of the way in which technological innovation and cost constraints often demand increased development effort.

It is important to reflect on the process that we have been following over the last few pages. The US Department of Energy recommends change analysis as a means of supplementing an initial ECF chart. The intention is to ensure that investigation consider a range of key events and the conditions that influence those events before any causal analysis is attempted. This approach is also recommended by the NASA guidelines for ‘Mishap Reporting, Investigating and Record-keeping’ [571] The Polar Lander case study illustrates a number of benefits that can be obtained from this complementary approach. In particular, the change analysis provides a good means of identifying the wider contextual issues that can often be overlooked by more event-based approaches. This is illustrated by the way in which change analysis helps to focus on the impact of managerial and organisational strategy. Our analysis has also indicated a number of potential weaknesses in the use of change analysis to inform the construction of ECF charts. Figure 10.18 only presents a small portion of the overall diagram. In ‘bespoke’ projects such as the Polar Orbiter mission, change analysis is likely to identify a vast range of potential differences from previous projects. It is important to reiterate that our case studies were deliberately chosen with this in mind, previous examples of ECF charts focus on the more routine analysis of incidents within the process industries [209].

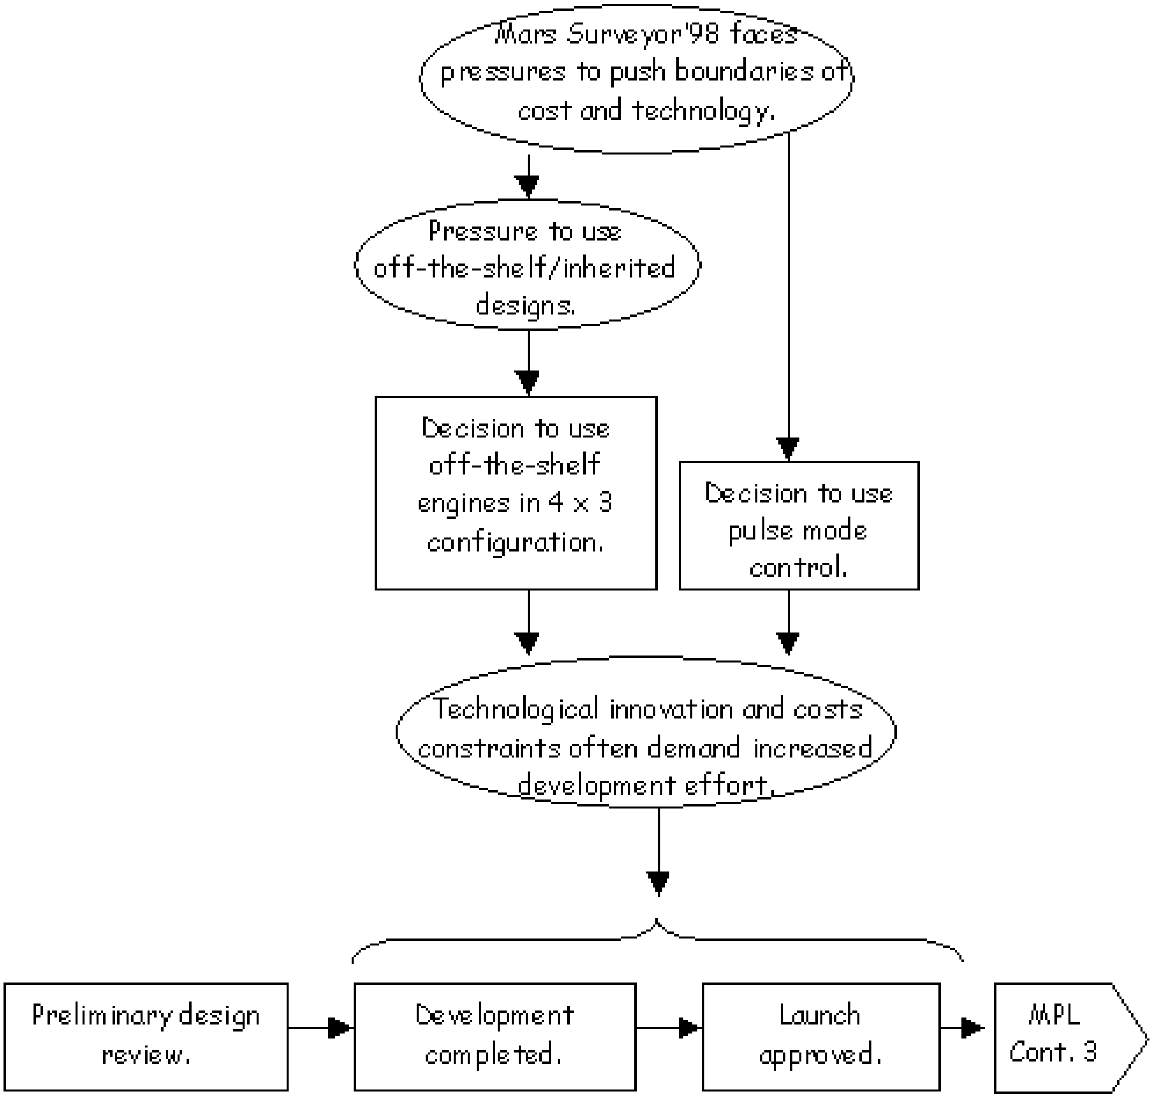
Mars Surveyor'98 faces pressures to push boundaries of cost and technology.
Pressure to use off-the-shelf/inherited designs.
Decision to use off-the-shelf engines in 4 x 3 configuration.
Decision to use pulse mode control.
Technological innovation and costs constraints often demand increased development effort.
Preliminary design review.
Development completed.
Launch approved.
MPL Cont. 3

Figure 10.18: Representing Technological Issues within an ECF chart (2)

Process: Changes in Development Practices and Reviews

Previous sections have identified differences between recommended risk management practices and the approach that characterised the Polar Lander's development. Many of the deficiencies can be explained by resource constraints. Others can be justified in terms of the practical challenges that such 'leading-edge' projects pose for current analysis techniques. The limited nature of the risk assessment process during the Polar Lander project did, however, have a number of knock-on effects. For example, previous NASA projects were typified by an extensive use of redundancy as a means of combating potential failures. The Shuttle's design was based on the maxim 'fail operational/fail operational/fail-safe'. One failure and the flight can continue but two failures and the flight must be aborted [565]. Even in these applications, however, it is not practical to develop fully redundant systems. In consequence, risk analysis guides the application of redundancy to the most mission-critical areas of a design. However, the lack of any system-wide hazard analysis arguably prevented the effective use of redundancy to protect against failure during key phases of the mission. It was noted that "certain MPL mission phases and sequences provide coverage only for parameter dispersions that conservatively represent stochastic dispersions, but unnecessarily fail to acceptably handle anomalously large parameter dispersions created by unmodeled errors or other non-stochastic sources" [570]. In particular, there was no functional backup if the Entry, Descent and Landing failed to follow an 'ideal' sequence of events. Table 10.8 summarises these knock-on effects that a limited risk analysis had upon the development of the Polar Lander mission.

Table 10.8 represents more general concerns about the models that guided the Lander's development. For instance, models were used to characterise the potential designs of the spacecraft as well as the environment in which it was intended to operate. Any inconsistencies, inaccuracies or omissions could have had profound consequences for the eventual success of the mission. Unfortunately, it is difficult to underestimate the complexity of constructing and validating such abstractions. Models that characterise one subsystem often influence, and are influenced by, many other subsystems. This creates considerable complexity because different aspects of a system are developed at different speeds. For example, thruster and software design lagged behind other Lander subsystems. Further problems complicated the use of predictive models. In particular, the small forces generated by the spacecraft could not be modeled to the level of accuracy that was required by the navigation plan. This called for precision navigation requirements that were incompatible with the spacecraft's design.

Validation and verification techniques can be used to test a potential design under simulated operating conditions. The results of such tests also provide insights into the utility of any models that guide systems development. Unfortunately, results can be compromised if validation tests are based on the same incorrect assumptions that guide mission development. Systems will perform well under simulated operating conditions that have little relationship with an eventual working environment. The problems of conducting such validation exercises are compounded by the managerial issues that complicate any multi-disciplinary development. Insufficient instrumentation, an error in the thermal model and poor communication between the propulsion and thermal groups produced inaccurate results from the Lander's thermal-vacuum tests. As a result, several design problems were not detected until after the launch. The Lander's validation "was potentially compromised in some areas when the tests employed to develop or validate the constituent models were not of an adequate fidelity level to ensure system robustness" [579].

NASA standards recommend independent verification and validation as a means of avoiding such problems [559]. Tests are conducted by organisations that are not involved in the development process. In consequence, they are less likely to follow the assumptions that are embodied within system models. External auditors may also be slightly more resilient to the internal pressures that complicate the conduct of integration tests within complex development teams. Unfortunately, this form of testing is expensive. On a resource-limited project, it must be focussed on those areas of a mission that are considered to be of prime importance. Technical difficulties further complicate the validation of complex systems. These problems prevented developers from testing system performance during the Entry, Descent and Landing phase under the Martian gravity of 3/8g. Partly as a result of this, the touchdown sensing software was not tested with the lander in the flight configuration and the software error was not discovered during the verification and validation program.

Figure 10.19 gathers together the products of the different forms of change analysis that have been conducted up to this point. These conditions describe the impact of changes in staffing policy and risk assessment practices. They also outline the effects of wider changes in NASA project management strategy and in development practices. These conditions collectively describe the context in which the Polar Lander was developed and launched. As more information becomes available about particular events, investigators can draw upon this contextual information to identify particular conditions that influenced those events. This approach provides a number of benefits. The conditions identified by change analysis need not be immediately associated with particular events. For example, conditions can emerged from the documents and statements that are gathered during a primary investigation. It can be difficult to identify particular events that are associated with the information that is provided by these documents. For instance, statistical comparisons of different levels of funding on various projects provide important information about the wider context in which an incident occurs. It would, of course, be possible to invent an event so that these conditions could be linked into an ECF

Prior/Ideal condition	Present Condition	Effects of change
Design is resilient beyond conservative stochastic parameter dispersions.	Design vulnerable to unmodeled errors or non-stochastic sources.	EDL Sequence fails under anomalous conditions
		No functional backup for several systems.
Spacecraft design should match mission requirements	Aspects of the design could not be modelled accurately enough for control	Small forces not accurately modelled for precision navigation.
Properly validated models should be used when testing is impossible	Some models not properly validated	Doubts over results for radar-terrain interaction.
		Doubts over dynamical control effects of pulse-mode propulsion.
Sufficient resources to assess interaction between propulsion, thermal and control subsystems	Thermal and software design lags behind other subsystems requiring these inputs. There was an error in the thermal model used to support thermal-vacuum tests. Insufficient instrumentation of the thermal-vacuum tests. Poor communication between propulsion and thermal groups.	Partial evaluation of propulsion, thermal and control interaction. Inadequate thermal-vacuum tests. Problem with catalyst bed heaters had to be handled prior to entry. Remaining concerns over uneven propellant drain from tanks during descent.
Sufficient resources to validate and verify software in landed configuration.	Flight software not subjected to 'system-level' tests.	Post-landing fault-response bugs only uncovered after mission loss.
		Touchdown sensing software untested with lander in flight configuration.

Table 10.8: Change Summary Table of MPL Process Issues.

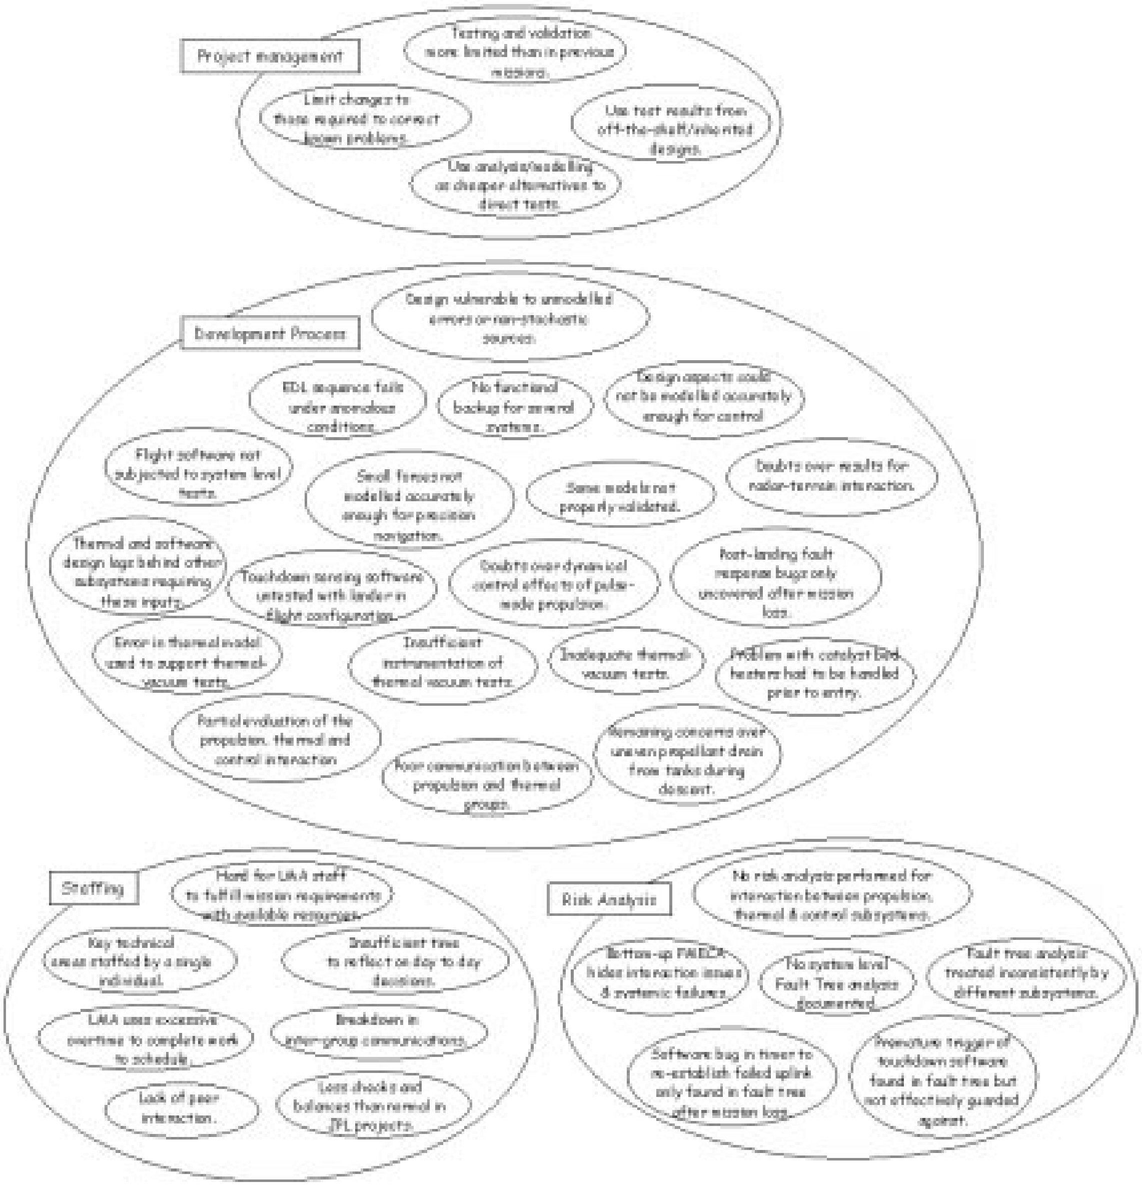
Project management
Testing and validation more limited than in previous missions.
Limit changes to those required to correct known problems.
Use test results from off-the-shelf/inherited designs.
Use analysis/modelling as cheaper alternatives to direct tests.
Development Process
Design vulnerable to unmodelled errors or non-stochastic sources.
EDL sequence fails under anomalous conditions.
No functional backup for several systems.
Design aspects could not be modelled accurately enough for control
Flight software not subjected to system level tests.
Small forces not modelled accurately enough for precision navigation.
Some models not properly validated.
Doubts over results for radar-terrain interaction.
Thermal and software design lags behind other subsystems requiring these inputs.
Touchdown sensing software untested with lander in flight configuration.
Doubts over dynamical control effects of pulse-mode propulsion.
Post-landing fault response bugs only uncovered after mission loss.
Error in thermal model used to support thermal-vacuum tests.
Insufficient instrumentation of thermal vacuum tests.
Inadequate thermal-vacuum tests.
Problem with catalyst bed heaters had to be handled prior to entry.
Partial evaluation of the propulsion, thermal and control interaction
Poor communication between propulsion and thermal groups.
Remaining concerns over uneven propellant drain from tanks during descent.
Staffing
Hard for LMA staff to fulfill mission requirements with available resources.
Key technical areas staffed by a single individual.
Insufficient time to reflect on day to day decisions.
LMA uses excessive overtime to complete work to schedule.
Breakdown in inter-group communications.
Lack of peer interaction.
Less checks and balances than normal in JPL projects.
Risk Analysis
No risk analysis performed for interaction between propulsion, thermal & control subsystems.
Bottom-up FMECA hides interaction issues & systemic failures.
No system level Fault Tree analysis documented.
Fault tree analysis treated inconsistently by different subsystems.
Software bug in timer to re-establish failed uplink only found in fault tree after mission loss.
Premature trigger of touchdown software found in fault tree but not effectively guarded against.

Figure 10.19: Using Change Analysis to Collate Contextual Conditions

chart. In contrast, Figure 10.19 shows how these contextual conditions can be gathered together for integration into an ECF chart, if and when investigators need to provide additional information about the conditions that affect particular events. Investigators are free to determine whether or not they should be explicitly associated with more detailed events. The complexity of ECF charts such as Figure 10.17 is an important consideration here. If all of the conditions represented in Figure 10.19 were explicitly linked to the different events that they influenced then the resulting ECF chart would rapidly become intractable. The task of determining the appropriate level of detail in such diagrams, therefore, forms an important component of the wider causal analysis.

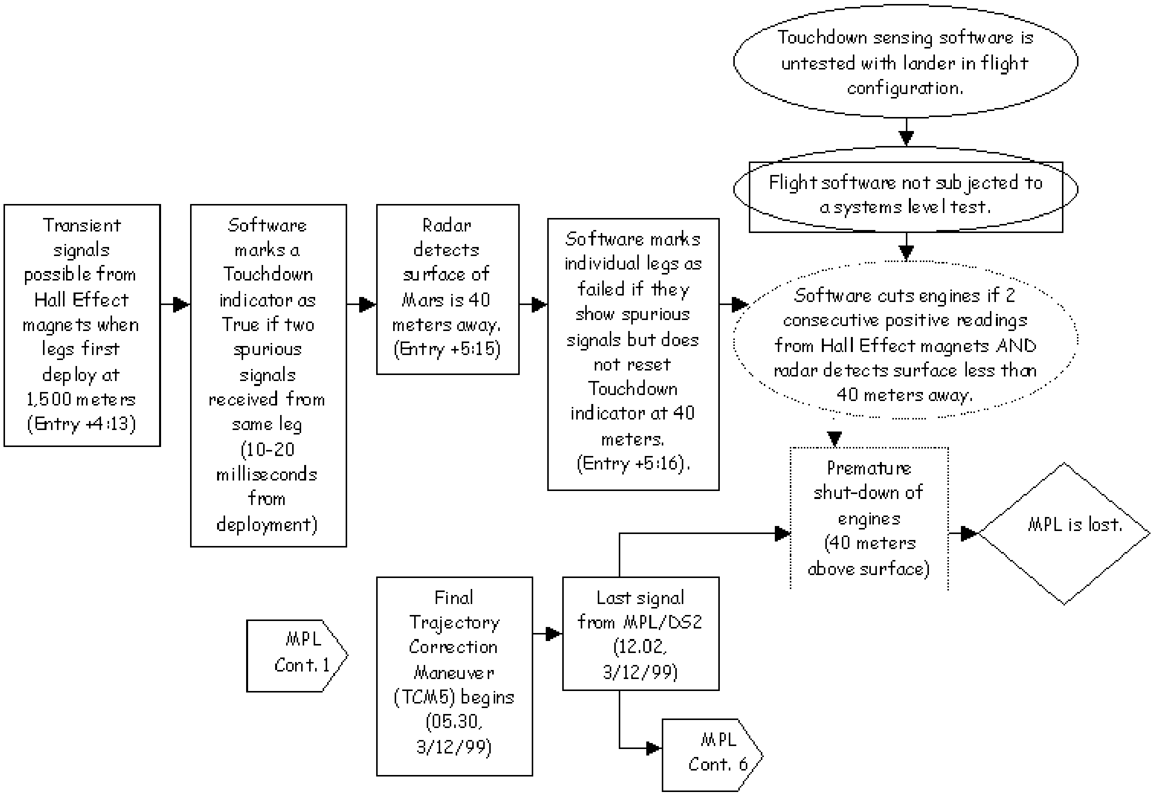
Touchdown sensing software is untested with lander in flight configuration.
Flight software not subjected to a systems level test.
Transient signals possible from Hall Effect magnets when legs first deploy at 1,500 meters (Entry +4:13)
Software marks a Touchdown indicator as True if two spurious signals received from same leg (10-20 milliseconds from deployment)
Radar detects surface of Mars is 40 meters away. (Entry +5:15)
Software marks individual legs as failed if they show spurious signals but does not reset Touchdown indicator at 40 meters. (Entry +5:16).
Software cuts engines if 2 consecutive positive readings from Hall Effect magnets AND radar detects surface less than 40 meters away.
Premature shut-down of engines (40 meters above surface)
MPL is lost.
MPL Cont. 1
Final Trajectory Correction Maneuver (TCM5) begins (05.30, 3/12/99)
Last signal from MPL/DS2 (12.02, 3/12/99)
MPL Cont. 6

Figure 10.20: Integrating Development Issues into an ECF chart (1)

Figure 10.20 illustrates how conditions can be introduced to provide further information about the events that are already represented within an initial ECF chart. In this case, the change analysis identifies that the **touchdown sensing software is untested with the lander in flight configuration**. It also identifies the more general point that the **flight software was not subjected to a systems level test**. These conditions both provide insights on the software problem that was identified in the Hall Effect sensors. This, in turn, led to the hypothesised failure scenario in which there was a premature shut-down of the lander's engines.

This analysis identifies a number of important caveats about our use of change analysis to drive the construction of ECF charts. In developing an initial ECF chart, we already identified the scenario in which the lander's engines were cut at forty meters above the planet surface. This helps to direct the subsequent analysis towards any changes that might have contributed to such a software failure. On the one hand, this can be seen as beneficial because it guides the allocation of finite investigatory resources. On the other hand, the generation of an initial hypotheses may bias any subsequent change analysis. This is especially important where there are considerable differences between each mission or run of a production process. Rather than considering the wider range of potential changes, analysts are biased towards those that support pre-existing hypotheses. This argument supports Mackie's ideas about causal fields that were introduced in Chapter 7 [508]. He goes on to develop the notion of a causal field that describes the normal state of affairs prior to any incident. Investigators try to identify the causes of an incident by looking for disturbances or anomalies within the causal field. This causal field is, therefore, a subjective frame of reference that individuals use when trying to explain what has happened in a particular situation. If a cause does not manifest itself within the causal field then its influence is unlikely to be detected. These ideas have a particular resonance in our use of change analysis. Both Table 10.19 and Figure 10.20 reflect subjective assumptions about what was 'normal' development practice. It was argued that sufficient resources should have been allocated to validate and verify software in landed configuration. Given that budgetary constraints affected almost every aspect of the Lander's development, the selection of this particular conditions provides insights not only about the incident itself but also about the investigator's causal field.

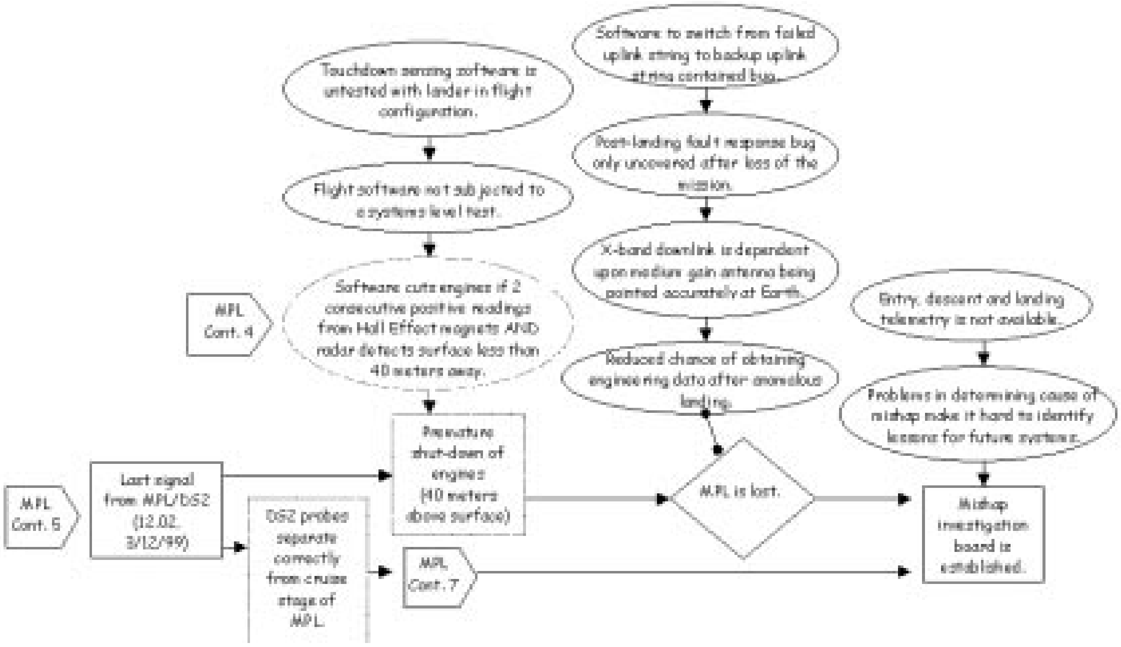
Touchdown sensing software is untested with lander in flight configuration.
Flight software not subjected to a systems level test.
Software cuts engines if 2 consecutive positive readings from Hall Effect magnets AND radar detects surface less than 40 meters away.
MPL Cont. 4
Premature shut-down of engines (40 meters above surface)
MPL Cont. 5
Last signal from MPL/DS2 (12.02, 3/12/99)
DS2 probes separate correctly from cruise stage of MPL.
MPL Cont. 7
Software to switch from failed uplink string to backup uplink string contained bug.
Post-landing fault response bug only uncovered after loss of the mission.
X-band downlink is dependent upon medium gain antenna being pointed accurately at Earth.
Reduced chance of obtaining engineering data after anomalous landing.
MPL is lost.
Entry, descent and landing telemetry is not available.
Problems in determining cause of mishap make it hard to identify lessons for future systems.
Mishap investigation board is established.

Figure 10.21: Integrating Development Issues into an ECF chart (2)

There is also a danger that the counterfactual arguments, which we have adopted, may also serve to compound the salience bias that we have described in the previous paragraph. Counterfactual reasoning encourages analysts to identify causes, which had they not occurred then the incident would not have occurred. There is a danger that this can lead to a search for 'silver bullets'; the minimal set of events that might have avoided the incident. This 'silver bullet' approach ignores Mackie's argument, introduced in Chapter 7 that there will be alternate 'causal complexes' that might lead to a future incident [508]. Mackie views a cause (in the singular) to be a non-redundant factor which forms part of a more elaborate causal complex. It is the conjunction of singular causes within the causal complex that leads to an outcome. The causal complex is sufficient for the result to occur but it is not necessary. There can be other causal complexes. By extension, the 'silver bullet' approach is likely to rectify singular causes within a causal complex. It is, however, likely to overlook other causal complexes that can lead to similar failures in the future. This is an abuse of counterfactual reasoning rather than a weakness of the approach itself. It is also important to distinguish between general and particular causation. A general cause is one which can be used to charaterise a number of different instances of the same factor. For example, poor situation awareness is a general cause of aviation accidents. In contrast, a particular cause is an instance of a general cause and describes a specific example of this more general problem. Hence we can have both general and particular, singular causes.

In the context of our analysis, there is a danger that change and barrier analysis might be used to support the preliminary hypotheses that are identified in ECF charts without examining the wider causal complexes identified by Mackie. Any subsequent root cause analysis will, therefore, be focussed on an extremely limited model of an incident. It is essential to stress noted that these dangers to not stem from the notations themselves. They are strongly related to the way in which those notations are used within particular incident investigations. In particular, the primary means of ensuring an adequate analysis of the causal complexes behind an incident is to expect the same level of review by peer investigators as one would expect during the design of any safety-critical system. Figure 10.21 illustrates how change analysis can be used to search for causal complexes beyond those that are identified in an initial ECF chart. This introduces conditions to denote that software to switch from a failed up-link string to a backup up-link string contained a bug and that post-landing fault response bug was only uncovered after the loss of the mission. As can be seen from the double headed edge in Figure 10.21 these conditions relate to problems in the communication system that could have contributed to the loss of the mission but not if the engines had indeed been cut at forty meters from the planet surface.

The previous paragraphs have argued that some of the software flaws were not detected because it was untested with the lander in flight configuration. There are both technical and financial barriers to such tests. NASA, therefore, advocates the use of formal reviews to supplement direct testing. These meetings are intended to increase consensus and confidence about a proposed design. For instance, the NASA Standard 5001 for the 'Structural design and test factors of safety for spaceflight hardware' states that:

> "Standard criteria cannot be specified for general use in designing structures for which no verification tests are planned. Projects which propose to use the no-test approach generally must use larger factors of safety and develop project-specific criteria and rationale for review and approval by the responsible NASA Center. For spacecraft and other payloads launched on the Space Shuttle, these criteria must also be approved by the Space Shuttle Payload Safety Review Panel prior to their implementation." [562]

Partly in response to the loss of the Climate Orbiter and the Polar Lander, NASA have recently published procedures for the 'Management of Government Safety and Mission Assurance Surveillance Functions for NASA Contracts' [568]. This identifies a continuum of oversight ranging from low intensity, periodic reviews to high intensity oversight, in which NASA managers have day-to-day involvement in the suppliers' decisionmaking processes. These different forms of oversight are coordinated through a surveillance plan that must be submitted within 30 days of any contract being accepted. The plan describes the safety and mission assurance functions that are necessary to assure that the contractor will meet project requirements. Independent agencies may be identified in this plan if they are to validate the results of any assurance functions. Surveillance plans must be revised to keep pace with changes in the contractors' operations. The plan and its revisions must be reviewed at least annually to determine whether or not it must be further revised. As mentioned, these requirements were not in place during the development of the Polar Lander. There are considerable dangers in applying standards that hold after an incident to identify deficiencies that led to any mishap. There, Table 10.9 restricts its analysis to those review activities that were recommended in documents such as [562] and [560].

The investigators found that the Polar Lander project did not have a documented review plan. It did, however, hold both formal and informal reviews. Each subsystem coordinated their own preliminary and critical design reviews. This informal approach was intended to reduce the level of bureaucracy that had been associated with assurance functions in other projects. This informal process was used to communicate concerns and generate requests for actions. Unfortunately, these subsystem reviews demonstrated varying levels of technical analysis. Some issues, such as the design of the G and H release nut, were examined in a meticulous and thorough manner. Others were not. For instance, the thermal control design interfaces were not mature enough to evaluate at propulsion systems critical design review. Had a subsequent review been scheduled then the developers might have discovered some the problems that were later experienced in flight.

Prior/Ideal condition	Present Condition	Effects of change
Subsystem Preliminary and Critical Design Reviews provide independent evaluation of key decisions	Contractors lacked necessary input from external sources	Flight System Manager chaired all subsystem reviews
		LMA staff approve closures on actions without independent technical support.
		Some actions did not adequately address concerns raised by reviews.

Table 10.9: Change Summary Table of MPL Review Issues.

A mission assurance manager tracked each review action to ensure that it was addressed by a written closure and that the closure was then approved by a relevant authority. This procedure was used to ensure that all actions and recommendations were closed prior to launch. These closures were, however, typically approved by LMA staff without any independent technical support. This need not have been a concern if some form of meta-level independent review had been conducted of these closures. As we have seen, however, budgetary constraints meant that there was minimal JPL technical support. LMA did not have their closures reviewed by Board members or by non-project LMA personnel. It was later argued that:

> "This limitation on technical penetration of the action items and their closure is not typical of JPL projects and was probably an unintended consequence of project funding limitations. Rather than following the typical process of choosing board chairpersons with technical expertise in functional areas from outside the project, the Flight System Manager was the chairperson of all the subsystem reviews." [579].

In passing, it is worth noting that the problems of developing effective assurance procedures for contracted work has been a recurring theme in recent NASA mishap reports [575]. This, in part, explains the subsequent development of a comprehensive set of standards and policies in this area.

Figure 10.22 provides a final illustration of the use of change analysis as a means of expanding an ECF chart. In this case, several further conditions are introduced to annotate the development and review events that have been identified by previous stages of the analysis. This figure again illustrates the problems of associating conditions with individual events. Parenthesis are again used below the event line to indicate the potential scope of these conditions. As with previous diagrams, it would be possible to refine the events shown in Figure 10.22 so that conditions can be more firmly rooted to particular moments during an incident. This is a subjective decision, I chose not to do it in this analysis because it would have forced me to invent a number of arbitrary events. The available evidence was not in a format where I could have such distinctions. In general, this reflects the difficulty of representing persistent constraints within event-based notations. Time-lines suffer from similar problems and the solutions were almost identical in Chapter 9. This remains an area of current research. For now, it is important to realise that our integration of change analysis and ECF charts has exposed a number of limitations in the application of this analysis technique for a complex, technological failure.

Previous sections focussed on the ways in which particular aspects of the Polar Lander's development may have contributed to the failure of this mission. In particular, we have identified instances in which this project adopted practices and procedures that differed from those advocated by senior management through published guidelines and policies. Limited funding and changes to

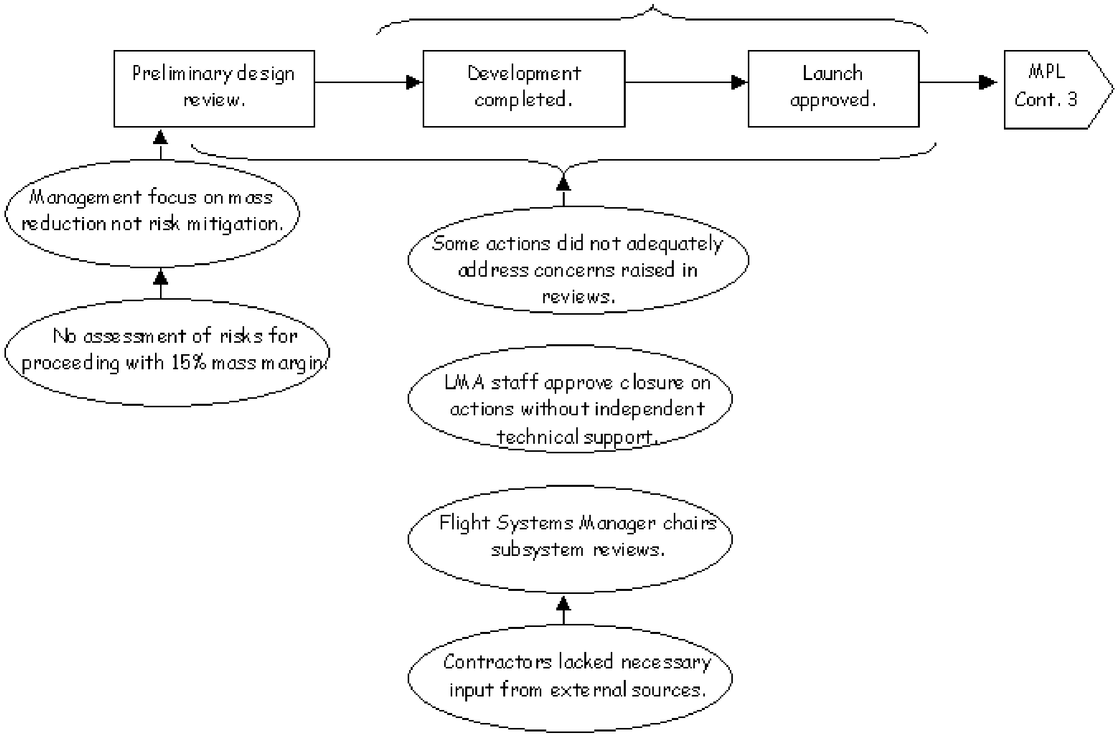
Preliminary design review.
Development completed.
Launch approved.
MPL Cont. 3
Management focus on mass reduction not risk mitigation.
No assessment of risks for proceeding with 15% mass margin
Some actions did not adequately address concerns raised in reviews.
LMA staff approve closure on actions without independent technical support.
Flight Systems Manager chairs subsystem reviews.
Contractors lacked necessary input from external sources.

Figure 10.22: Integrating Review Issues into an ECF chart

NASA's subcontracting practices helped to place heavy burdens upon the available staff. These burdens, together with particular skill shortages, had an adverse effect on the risk assessments that are intended to guide subsequent development. As a result, a number of technical decisions were made that could not easily be justified in retrospect. For example, the lack of telemetry during the Entry, Descent and Landing phase created considerable problems for investigators who must feed any relevant lessons into current and future projects. Furher problems arose from the technical and financial barriers that prevented development teams from testing all aspects of the Polar Lander's design. Such tests might have helped to identify potential problems that were not identified during a hazard analysis. Instead, a number of problems were discovered after the craft was in flight. Such problems also illustrate the way in which the Polar Lander's project reviews had failed in their meta-level role of assuring mission success.

It is important to stress that the previous tables have been guided by an implicit form of change analysis that is apparent in the documents and records that were produced by the NASA investigators. In order to identify potential shortcomings that might have affected the mishap, they first had to analyse the recommended practices for similar development projects:

> "NASA currently has a significant infrastructure of processes and requirements in place to enable robust program and project management, beginning with the capstone document: NASA Procedures and Guidelines 7120.5. To illustrate the sheer volume of these processes and requirements, a partial listing is provided in Appendix D. Many of these clearly have a direct bearing on mission success. This Boards review of recent project failures and successes raises questions concerning the implementation and adequacy of existing processes and requirements. If NASA's programs and projects had implemented these processes in a disciplined manner, we might not have had the number of mission failures that have occurred in the recent past." [579]

For example, the software component of the Lander development was covered by NASA standard NASA-STD-2100-91 (Software Documentation, [558]), by NASA-STD-2201-93 (Software Assurance,

[559]), by NASA-STD-2202-93 (Software Formal Inspections, [560]) and by a draft form of NASA-STD-8719.13A (Software Safety, [563]). This illustrates an important limitation of change analysis. In an organisation as complex as NASA, it is likely that there will be a significant body of information about recommended practices. It can be difficult or impossible for any individual to continually assess whether their project conforms to all of the available guidelines. As a result, it is likely that most projects will differ from the ideal. It can also be difficult for developers to learn more about successful practices from other projects. One means of addressing this problem is to provide developers with means of searching for appropriate guidelines and lessons learned. NASA provide a web-based interface to their standards library for this purpose. By extension, it can also be argued that same facilities ought to be available to help inspectors search for incidents in which these standards were not followed. Such tools can be used to identify emerging patterns of related failures within a database of incidents. Chapter 15 will describe some of these systems in more detail. In contrast, the following section goes on to show how ECF charts can be used to direct a causal analysis of the Polar Lander and Climate Orbiter case studies.

10.3 Stage 2: Causal Analysis

This section goes on to describe how a number of analytic techniques can be used to distinguish causal events from the mass of contextual events and conditions that are identified in preliminary ECF charts. In particular, ECF Analysis, Tier Diagramming and Non-compliance Analysis are used to filter the mass of information that is gathered during primary and secondary investigations.

10.3.1 Causal Factors Analysis

The Department of Energy guidelines argue that ECF charting must be conducted to a sufficient level of detail and that this depends upon *both* change *and* barrier analysis [207]. The NASA guidelines, NPG 8621.1, are ambiguous in this respect [571]. Barrier analysis appears as an item in the Mishap Board Checklist (Appendix J-3) but not in the list of recommended investigation techniques where guidance is provided on the other two complementary approaches. Irrespective of whether both analytical techniques are used to derive an ECF chart, the next stage is to analyse the resulting diagram to identify the causes of an incident. This, typically, begins with the event that immediately precedes the incident. The Department of Energy guidelines suggest that investigators must ask would the incident have occurred without this event?. If the answer is yes then the analyst progresses to the next event; the event is assumed not to have had a significant impact on the course of the incident. However, if the answer is no then a number of further questions must be asked about the both the event and the conditions that are associated with it. This illustrates how causal factor analysis relies upon counterfactual argument.

A number of problems complicate this first stage of the analytical method. The first issue centres on the relationship between events and conditions. Previous sections have argued that conditions “(a) describe states or circumstances rather than happenings or occurrences and (b) are passive rather than active” [209]. Problems arise when a condition is associated with an event that is not considered to be central to the causes of an incident, i.e., the answer to the previous counterfactual question is yes . For instance, it might be argued that the Climate Orbiter might still have been lost even if more staff had transitioned from development to operations. In this case, investigators might then neglect the effect of the associated condition that the Mars Climate Orbiter is the first project for the multi-mission Mars Surveyor Operations project. It can be argued that such conditions are irrelevant because they do not directly affect the counterfactual argument that drives causal factor analysis. It can also be argued that this form of analysis places unnecessary importance on specific events and that it neglects the context in which an incident occurs. Such caveats are important because many event-based modelling techniques force investigators to invent ‘arbitrary’ events so that they can represent important elements of this context. For example, failures of omission have to be represented as negative events within an ECF line. This provides investigators with the only means of representing the conditions that influenced the omission. For example, the decision not to perform TCM-5 was influences by the failure to understand the significance of the AMD data. This, in turn, was influenced by conditions that ranged from management changes through to a reliance on Doppler shift and the Deep Space network for tracking data. This example clearly illustrates that it is the conditions that are more important for future safety than the 'non-event'.

Causal factor analysis is further complicated by the difficulties of applying counterfactual reasoning to complex, technological failures. For instance, how can we be sure that the Climate Orbiter would have succeeded if the Small Forces bug had been counteracted by TCM-5? There might have been other unidentified problems in the navigation software. Alternatively, TCM-5 might itself have introduced further problems. The key point here is that the previous counterfactual question refers to a particular incident. It does not ask 'would any incident would have occurred without this event?'. Investigators cannot, typically, provide such general guarantees.

Further complications arise from multiple independent failures. These occur when an investigation reveals two or more problems that might have led to an incident. Multiple independent failures are denoted on ECF charts by different chains of events and conditions that lead to the same incident symbol. Our analysis of the Polar Lander identified two of these chains. One leads from the failure of the touchdown sensing logic. The other represents problems in the communications systems. These independent failures create problems for counterfactual arguments because the incident might still have occurred if either one of them was avoided. An investigator would answer 'yes' to the question 'would the incident have occurred without the Hall Effect sensor problem?'. Conversely, they could also answer 'yes' to the question 'would the incident have occurred without the communications problems after landing'. According to the ECF method they would then disregard these events and continue the analysis elsewhere! This problem can be avoided if investigators construct and maintain multiple ECF charts to represent each of these different paths. This approach has some drawbacks. For instance, it can be argued that similar events led to the touch-down sensing bugs and the software problems in the communications up-link. These common causes would then be artificially separated onto different ECF charts in order to preserve the method, described above. An alternative means of avoiding this problem is to require that investigators repeat the counterfactual question for each path that leads to an incident symbol. The question then becomes 'would the incident have occurred in the manner described by this ECF path without this event?'.

The complex issues surrounding counterfactual reasoning about alternative hypotheses does not simply affect the Polar Lander and Climate Orbited case studies. It is a research area in its own right. Byrne has conducted a number of preliminary studies that investigate the particular effects that characterise individual reasoning with counterfactuals [123, 124]. This work argues that deductions from counterfactual conditionals differ systematically from factual conditionals and that, by extension, deductions from counterfactual disjunctions differs systematically from factual disjunctions. This is best explained by an example. The statement that 'the Climate Orbiter either re-entered heliocentric space or impacted with the surface' is a factual disjunction. Byrne argues that such sentences impose additional burdens on the reader if they are to understand exactly what happened to the Climate Orbiter. In the general case, they must also determine whether both of the possible outcomes could have occurred. The statement that 'the Climate orbiter would have re-entered heliocentric space or would have impacted with the surface' is a counterfactual disjunction. Byrne argues that this use of the subjunctive mood not only communicates information about the possible outcome of the mission but also a presupposition that neither of these events actually took place. There has, to date, been no research to determine whether these insights from cognitive psychology can be used to explain some of the difficulties that investigators often express when attempting to construct complex counterfactual arguments about alternative scenarios. In particular, the use of counterfactual disjunctions in our analysis of the Polar Lander is specifically not intended to imply that neither actually took place. It, therefore, provides a counter-example to Byrne's study of the everyday use of this form of argument.

Figure 10.23 presents an excerpt from the ECF chart that represents the failure of the Polar Lander mission. As can be seen, this diagram focuses on the events and conditions that may have contributed to the loss of the Deep Space 2 probes. The following paragraphs use Figure 10.23 to illustrate the application of the analytical techniques described above. In contrast to the Climate Orbiter and the Lander itself, we have not applied change or barrier analysis to this portion of the initial ECF chart. The decision to focus on this aspect of the incident is entirely intentional.

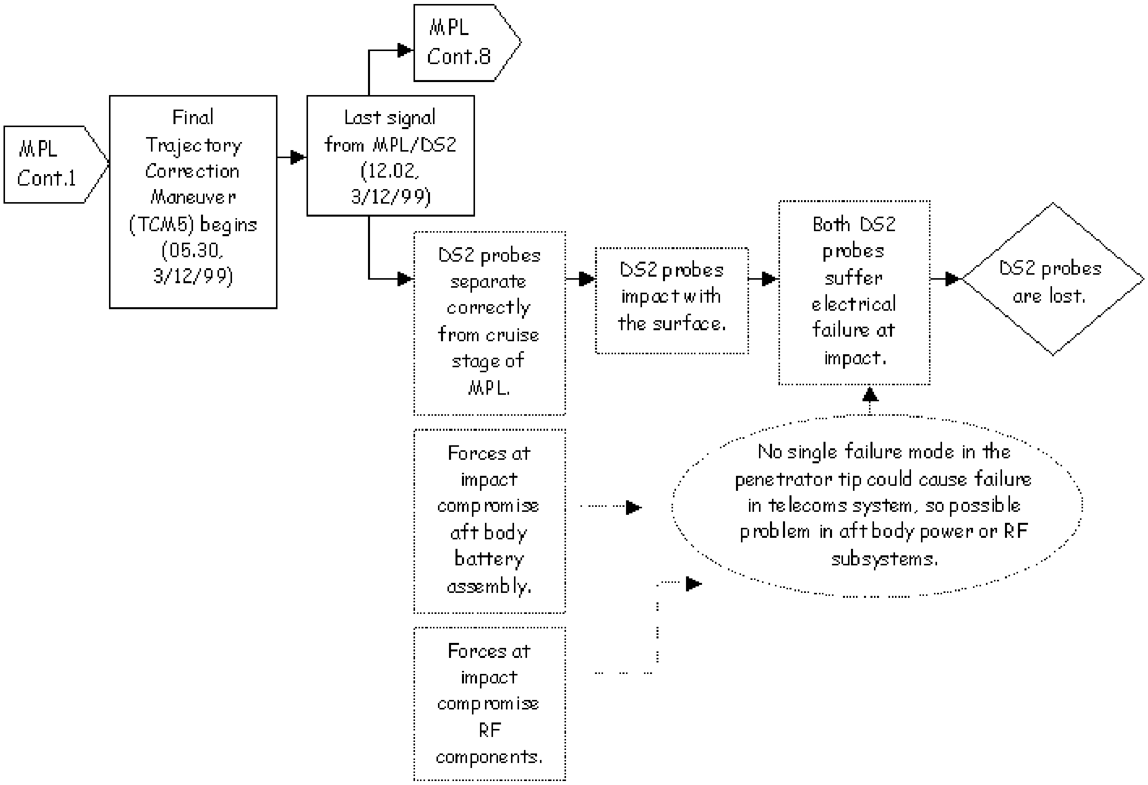
MPL Cont.8
MPL Cont.1
Final Trajectory Correction Maneuver (TCM5) begins (05.30, 3/12/99)
Last signal from MPL/DS2 (12.02, 3/12/99)
DS2 probes separate correctly from cruise stage of MPL.
DS2 probes impact with the surface.
Both DS2 probes suffer electrical failure at impact.
DS2 probes are lost.
Forces at impact compromise aft body battery assembly.
No single failure mode in the penetrator tip could cause failure in telecoms system, so possible problem in aft body power or RF subsystems.
Forces at impact compromise RF components.

Figure 10.23: An ECF chart of the Deep Space 2 Mission Failure

The subsequent paragraphs show how causal factor analysis can be used to check whether change and barrier analysis has identified the precursors and conditions that affect the potential causes of failure. As mentioned, causal factor analysis begins with the event that immediately precedes the incident symbol. Previous paragraphs have argued that the answer to this question is bounded by the particular ECF path that is being considered. It would, therefore, be necessary to repeat the analysis for each alternate paths leading to the same incident. Fortunately, Figure 10.23 shows a single event chain leading to the accident.

The investigator must ask whether the failure would have occurred if it was not the case that both of the DS2 probes suffer electrical failure at impact? If the answer were yes, the incident could have occurred without this failure, then the event can be classified as a contextual detail. The analysis would then move on to preceding events. In this case, however, if the electrical failure had not occurred then the probes would not have been lost. If we had omitted this event from our model, we would not have had a coherent explanation of the failure. This counterfactual argument suggests that this event is a contributory factor and that further causal factor analysis should be conducted. This causal factor analysis is based around a number of questions that are intended to ensure that analysts have identified sufficient information about key events. This information is necessary to drive any subsequent root cause analysis. It is important to stress, however, that many of the details that emerge from a causal factor analysis may already have been identified during previous stages of barrier and change analysis. This penultimate stage, therefore, provides additional assurance in the results of these other analytical techniques. The US Department of Energy guidelines argue that investigators must review the results of this analysis so that 'nothing is overlooked and that consensus has been achieved' [207].

Table 10.10 records the results of an initial causal factor analysis for the electrical failure event that precedes the loss of the probes shown in Figure 10.23. As can be seen, the intention behind the questions that drive the causal factor analysis is to expand on the summaries that label the

Event : Both DS2 Probes Suffer Electrical Failure at Impact	
What led to the event?	There was not enough time to conduct an impact test with a complete probe in flight configuration. Cost constraints and technical barriers also prevented such a validation.
What went wrong?	1. There was no system-level impact test of a flight-like RF subsystem. Mechanical and structural validation took place at the level of brassboard and breadboard components. Many components were not electronically functional. This limited pre-test and post-test DC continuity checks. 2. The flight battery cell lot was delivered too late to be impact tested. Validation arguments were based on a preceding lot of 8 identical cells. However, one of these was physically damage during a test but did not fail catastrophically.
How did the barriers fail?	The program exploited non-destructive tests and analytical modelling whenever possible. This was in-line with the objectives of the Faster, Better, Cheaper strategy. However, analytical models of high g impacts are unreliable and so flight qualification should have been demonstrated by tests on representative samples of flight hardware.
Who was involved in the event?	Two peer review meetings and three project level reviews established “proceed to launch” concurrence from JPL and NASA upper management. If the project team had forced an impact test for the RF subsystem and the fully integrated, powered probe then they might have missed the launch.
Is the event linked to a more general deficiency?	Many events and conditions in the Polar Lander’s ECF charts that relate to validation and review problems. The Faster, Better, Cheaper strategy is relevant to different events and conditions also.

Table 10.10: ECF Analysis of the Deep Space 2 Failure.

ECF chart. The ECF chart is used to show *when* an event occurred. The causal factor analysis expands this to capture *what* went wrong, *why* barriers failed and *who* was involved in the event. It should be noted that these questions are a subset of those proposed by the US Department of Energy [207]. This is intended to simplify the causal factor analysis and broaden its application to include the complex, technological failures that are addressed in this chapter. It should also be noted, however, that these questions can be amended to reflect the insights that are gained during subsequent investigations. For instance, we initially had replaced *who was involved in the event?* with the question *who was responsible for the barrier?*. This original version was removed after some investigators used the answer to directly assign blame for the incident even though barriers may have been breached by a pathological conjunction of environmental behaviours and system failures.

As can be seen, the causal factor analysis in Table 10.10 helps to collate information about the development of the probes. It describes how the flight cell battery lot was delivered too late to be impact tested. Table 10.10 also includes information about validation activities. There was insufficient time to conduct a powered, fully integrated impact test on the probe communications system. Finally, it identifies groups who were responsible in approving the "proceed to launch" decision in spite of these potential concerns. These observations were not explicitly identified during previous stages in the generation of the ECF chart. They, therefore, can be interpreted as omissions that are exposed by the explicit questions in the form shown in Table 10.10. Additional events can be introduced into Figure 10.23 to represent these insights prior to the eventual root cause analysis.

The final question in Table 10.10 looks beyond the specific event that forms the focus of this analysis. In particular, it prompts the investigator to identify whether or not a particular failure forms part of a wider pattern. It follows that such annotations are likely to be revised as the causal factor analysis is repeated for many different events in an ECF chart; patters may only emerge during the subsequent analysis. This question also provides an opportunity to explicitly identify any similarities with previous events during other incidents. Subsequent chapters will describe tools and techniques that can be used to identify common features amongst a number of different incidents. For now, however, it is sufficient to observe that primary and secondary investigations often uncover superficial similarities between the events that contribute to different incidents. These potential similarities must be investigated to determine whether or not different incidents do indeed begin to form a pattern of failure.

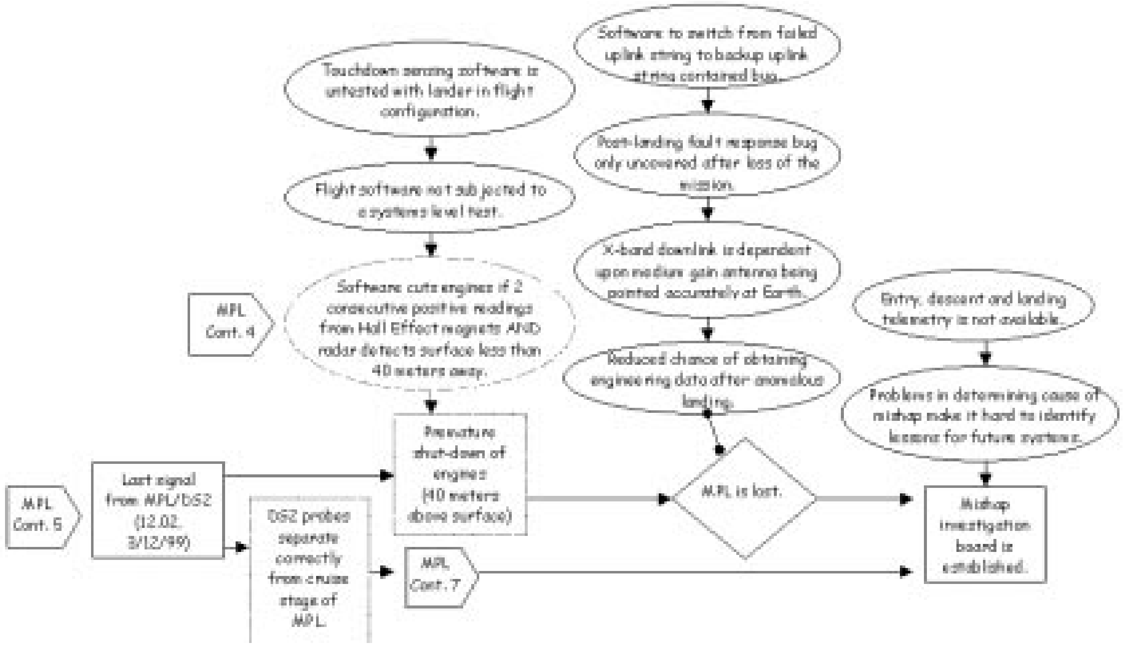
Touchdown sensing software is untested with lander in flight configuration.
Flight software not subjected to a systems level test.
Software cuts engines if 2 consecutive positive readings from Hall Effect magnets AND radar detects surface less than 40 meters away.
MPL Cont. 4
Software to switch from failed uplink string to backup uplink string contained bug.
Post-landing fault response bug only uncovered after loss of the mission.
X-band downlink is dependent upon medium gain antenna being pointed accurately at Earth.
Reduced chance of obtaining engineering data after anomalous landing.
Entry, descent and landing telemetry is not available.
Problems in determining cause of mishap make it hard to identify lessons for future systems.
MPL Cont. 5
Last signal from MPL/DS2 (12.02, 3/12/99)
Premature shut-down of engines (40 meters above surface)
DS2 probes separate correctly from cruise stage of MPL.
MPL Cont. 7
MPL is lost.
Mishap investigation board is established.

Figure 10.24: An ECF chart of the Polar Lander Mission Failure

The causal factor analysis in Table 10.10 is untypical because we have not presented any previous barrier or change analysis to identify further events and conditions leading to the loss of the Deep Space 2 mission. This was intentional because some investigations may not have the necessary resources to conduct these intermediate forms of analysis. As we have seen, it is possible to move straight from a high-level preliminary ECF chart such as Figure 10.4 to the analysis in Table 10.10. For higher consequence failures, such as the Mars Global Surveyor missions, it is likely that any causal factor analysis will build upon barrier and change analysis. Figure 10.24, therefore, integrates the events and conditions that were identified in the previous analysis of the Polar Lander incident. The relative complexity of this figure, even with the use of continuation symbols, indicates the complexity of the incident. It also provides an overview of the investigations that precede causal factor analysis.

The incident symbol in Figure 10.24 is preceded by an event, labelled Premature shut-down of engines (40 meters above the surface), and by a condition, labelled Reduced chance of obtaining engineering data after anomalous landing. Previous sections have, however, explained that these events are mutually exclusive. This is denoted by the double-headed link between the condition and the incident symbol. If the engines had been shut-down at 40 meters then the Lander would have been destroyed on impact with the planet surface. In consequence, any problems with the communications systems are unlikely to have had a significant impact on the loss of the mission. There is a very small probability that it could have survived such an event but the NASA investigation team did not consider that it was worth pursuing. In consequence, the causal factor analysis focuses on the event that is associated with the engine shut-down.

Causal factor analysis begins by asking whether the failure would have occurred if there had not been *premature shut-down of engines (40 meters above the surface)*. The answer to this question is assumed to be no. This is the only event in the ECF chart of Figure 10.24 that leads to the loss of the mission. The enquiry process, therefore, follows the same pattern as that established for the loss of the Deep Space 2 probes. Table 10.11 summarises the answers to the questions that drive the causal factor analysis.

Table 10.10 was derived without any intermediate barrier or change analysis. In contrast, Table 10.11 benefits from the more sustained analysis described in previous sections. In consequence, the ECF prompts may simply reiterate information that was identified by the earlier forms of analysis. The premature shut-down stemmed from a spurious touchdown signal from the Hall Effect sensors. The software did not reset a variable that was set in response to this spurious signal and this ultimately indicated that the Lander had contact with the surface when it was still some 40 meters from touch-down. It is, however, likely that the causal factor analysis will prompt some novel observations. For example, Table 10.11 briefly explains how the developers were keen to balance the loading on processors during the Entry, Descent and Landing phase. This contributed to the software failure because processors sampled the Hall Effect sensors well before reaching 40 meters. The intention was to avoid any sudden processing peaks that might have been incurred by starting to poll these devices at the point at which their input was needed.

The causal factor analysis also poses some questions that were not directly addressed during previous stages in the investigation. The change analysis of the Polar Lander failure did not explicitly address the reasons *why* particular barriers failed to detect the potential bug in the landing software. As can be seen from Table 10.11, the XB0114 requirements document did not explicitly consider the possible failure modes for the landing logic. The software engineers were not informed of the possibility of transient signals when the legs first deployed. The need to guard against such spurious signals was not explicitly included within the the Software Requirements Specification. In consequence, this requirement was not propagated into subsequent test protocols..

Table 10.11 illustrates further benefits of this analysis technique. ECF charts, typically, stretch over many pages. As can be seen from Figure 10.24, this can separate key events during the analysis and testing of a system from the point at which it is presumed to fail. The drafting of XB0114 occurred long before contact was lost with the Polar Lander. ECF charts, such as that shown in Table 10.11, help to trace the impact that distal events and conditions have upon catalytic failures. This is a significant benefit for complex, technological incidents. For example, our analysis of the Polar Lander failure and the associated loss of the Deep Space 2 probes extends to well over fifty nodes. This analysis is still at a relatively high level of abstraction. Several other investigations have

Event : Premature Shut-down of engines	
What led to the event?	Software did not reset a variable to denote that a spurious touchdown signal had been detected. This variable was read when the touchdown sequence was enabled at forty meters. The lander had an approximate velocity of 13 meters per second, in Martian gravity this accelerates to 22 meters per second at impact.
What went wrong?	Data from the engineering development deployment tests, flight unit deployment tests and Mars 2001 deployment tests showed a spurious reading in the Hall Effect touchdown sensor during landing leg deployment. These spurious signals can continue long enough to be detected as valid. Software that was intended to protect against this did not achieve the intended result. Spurious signals were retained until the sensing logic was enabled at 40 meters from the surface.
How did the barriers fail?	Requirements document (XB0114) did not explicitly state possible failure modes. Software engineers were not told about the transient failures. The system level requirements included a clause that might have alerted engineers to this problem but it was not included in Software Requirements Specification. The transient protection requirement was not, therefore, tested in either the unit or system level tests nor was it looked for in software walk-throughs. There was also an attempt to load balance on the processor so sampling started well before the 40 meter threshold. Product Integrity Engineer for Hall Effect sensors was not present at walk-throughs.
Who was involved in the event?	Software engineers, Product Integrity Engineers.
Is the event linked to a more general deficiency?	Problems in the Polar Lander software for the communications up-link. Software problems also affected Climate Orbiter and Stardust.

Table 10.11: ECF Analysis of the Polar Lander Failure.

produced ECF charts that contain over one thousand events and conditions. In such circumstances, it is essential that analysts have some means of summarising and collating information about the key events that contribute to an incident.

Previous paragraphs have used causal factor analysis to drive a more detailed consideration of the events that immediately precede the loss of the Polar Lander and the Deep Space 2 mission. If there was sufficient funding, then investigators would continue the analysis for each events on every path to the incident. If the incident would not have occurred without this event then the supplementary questions in Tables 10.10 and 10.11 would be posed. This approach might be seen to impose unwarranted burdens upon an investigation team. As we have seen, however, it can help to identify new insights into the events leading to high-criticality failures even if other forms of analysis have already been applied. Brevity prevents an exhaustive exposition of this approach. In contrast, Figure 10.25, therefore, presents an ECF chart for the loss of the Climate Orbiter. As can be seen, this diagram integrates the events and conditions from several previous diagrams. These earlier figures included continuation symbols. Figure 10.25 uses these to piece together a more complete view of the incident. As before, however, it is not possible to provide a single legible diagram of all of the events and conditions that were identified by the previous use of change and barrier analysis.

One of the reasons for focusing on Figure 10.25, rather than repeating the causal factor analysis of Deep Space 2 or the Polar Lander, is that it can be used to illustrate the distinction between contextual and causal factors. As before, the analysis starts from the event that precedes the incident. In this case, we must consider whether the incident would still have occurred if the Last signal from MCO (09:04:52, 23/9/99) had not occurred. It seems clear that the incident might still have occurred even if this event had not taken place. If we had omitted this event from our model, we would still have had a coherent explanation of the failure. It, therefore, represents a contextual rather than a causal factor. It is an event that helps our understanding of the incident but it is not necessary to our view of the incident. The analysis, therefore, moves to the event that immediate precedes the previous focus for the analysis. In this case, we must consider whether the incident would have occurred if the Mars Orbital Insertion had not taken place. Again, this event can be omitted without jeopardising the account of the failure. Similarly, the end of the cruise phase is not necessary to a causal explanation of the loss of the Climate Orbiter. The analysis, therefore, moves to the event labelled TCM-5 is discussed but not executed (16-23/9/99).

This event illustrates the complexity of counterfactual reasoning if investigators are not careful about the phrases that are used to label the nodes in an ECF chart. They must determine if the incident would have occurred if it was not the case that TCM-5 is discussed but not executed. The complexity in answering this question stems in part from a mistake in the construction of the ECF chart. As mentioned previously, events should be atomic statements. The previous label refers to both the discussion of the maneuver and to the decision not to implement it. In consequence, Figure 10.25 can be simplified by re-writing this event as It is decided not to execute TCM-5. The discussions surrounding this decision could be shown as an additional, secondary chain of events. It would have been easy to write this chapter with the 'correct' version from the start. This was not done because it is important to emphasise that the development of an ECF chart is an iterative process. It does not guarantee the construction of an 'error free' diagram. In consequence, causal factor analysis provides important checks and balances that can be used to support any causal investigation.

The counterfactual question based on the re-writing of the event now becomes would the incident would have occurred if it was not the case that it was decided not to execute TCM-5? This is equivalent to would the incident would have occurred if it was decided to execute TCM-5? Using the counterfactual question as a test, this event can be considered to have contributed to the failure. The incident need not have occurred if TCM-5 had been executed. A number of caveats can be raised to this argument. For instance, this assumes that that TCM-5 would have been performed correctly. It also assumes that the decision would have been taken when it was still possible to correct the trajectory of the Climate Orbiter prior to insertion. There are further complexities. If we ask the subsidiary question would the ECF chart still represent a plausible path to the incident without the event then it can be argued that the omission of TCM-5 did not cause the incident. It provided a hypothetical means of getting the system back into a safe state. It is, therefore, qualitatively different from the active

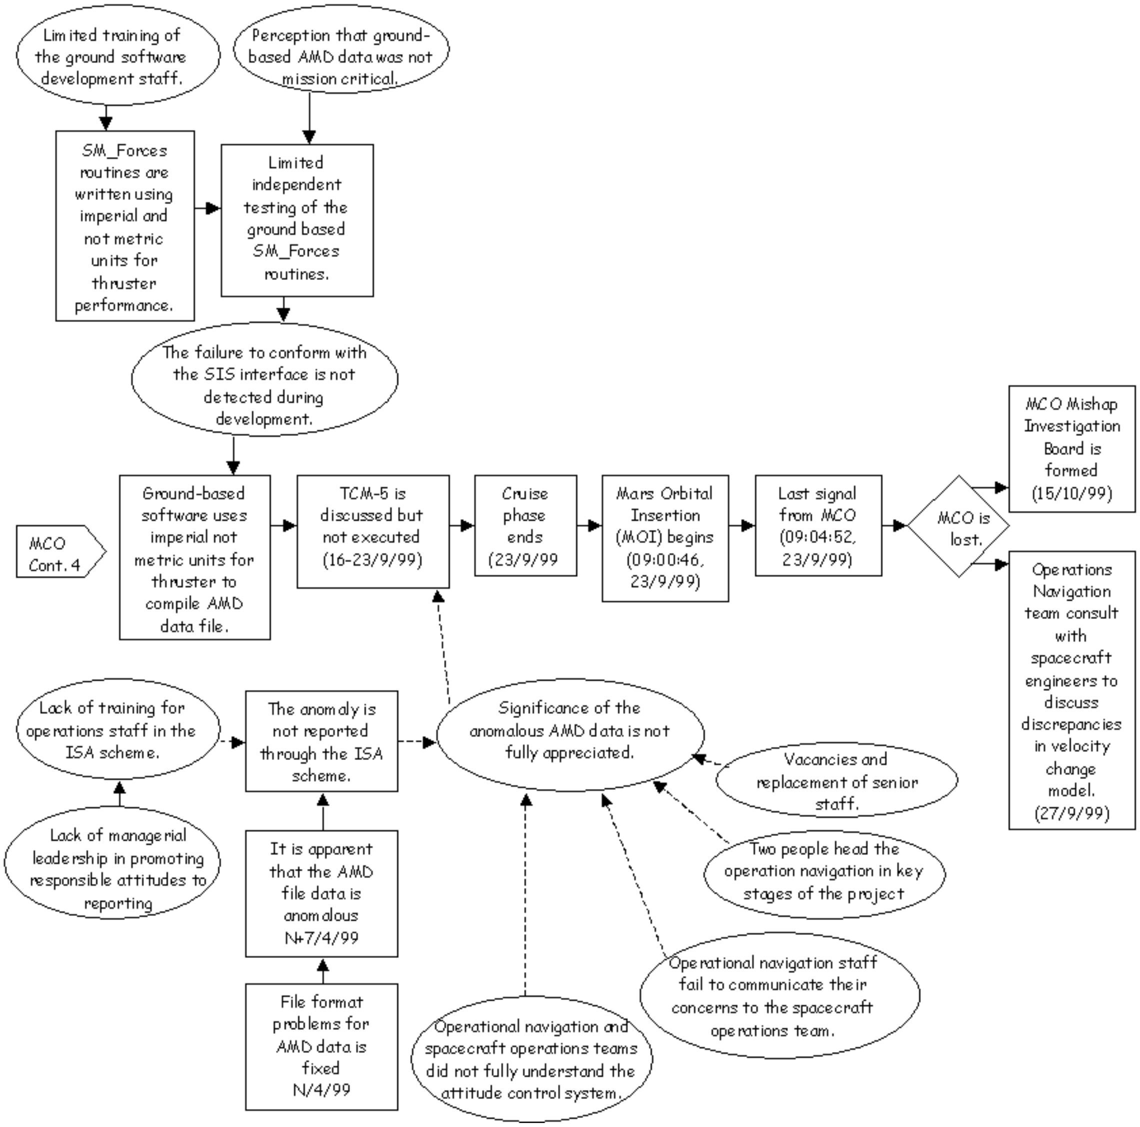
Limited training of the ground software development staff.
Perception that ground-based AMD data was not mission critical.
SM_Forces routines are written using imperial and not metric units for thruster performance.
Limited independent testing of the ground based SM_Forces routines.
The failure to conform with the SIS interface is not detected during development.
MCO Mishap Investigation Board is formed (15/10/99)
MCO Cont. 4
Ground-based software uses imperial not metric units for thruster to compile AMD data file.
TCM-5 is discussed but not executed (16-23/9/99)
Cruise phase ends (23/9/99
Mars Orbital Insertion (MOI) begins (09:00:46, 23/9/99)
Last signal from MCO (09:04:52, 23/9/99)
MCO is lost.
Operations Navigation team consult with spacecraft engineers to discuss discrepancies in velocity change model. (27/9/99)
Lack of training for operations staff in the ISA scheme.
The anomaly is not reported through the ISA scheme.
Significance of the anomalous AMD data is not fully appreciated.
Vacancies and replacement of senior staff.
Lack of managerial leadership in promoting responsible attitudes to reporting
It is apparent that the AMD file data is anomalous N+7/4/99
Two people head the operation navigation in key stages of the project
File format problems for AMD data is fixed N/4/99
Operational navigation staff fail to communicate their concerns to the spacecraft operations team.
Operational navigation and spacecraft operations teams did not fully understand the attitude control system.

Figure 10.25: An ECF chart of the Climate Orbiter Mission Failure

failures that are addressed in previous paragraphs.

The previous paragraph has argued that TCM-5 is a causal event according to the strict application of our counterfactual argument. We have, however, also identified counter arguments. The omission of TCM-5 was not a causal event because even if the decision had been taken to perform this operation there is no guarantee that it would have prevented the incident from occurring. This ambiguity stems from the difficulty of counterfactual reasoning about contingent futures. Not only do we have to imaging the there was a decision to implement TCM-5 but we also have to be sure that it would have avoided the incident. The complexity of such arguments has led a number of research teams to apply mathematical models of causation to support informal reasoning in accident investigation [470, 118]. These models attempt to provide unambiguous definitions of what does and what does not constitute a causal relation. They are, typically, based on a notion of distance between what actually happened and what might have happened under counterfactual arguments. A scenario in which TCM-5 was performed and did avoid the incident might be argued to be too far away from the evidence that we have about the actual incident. Such approaches offer considerable benefits; they can be used to prove that different investigators exploit a consistent approach to incident analysis. Unfortunately, the underlying formalisms tend to be unwieldy and error-prone especially for individuals who lack the appropriate mathematical training. A related point is that mathematical definitions of causation are frequently attacked because they fail to capture the richness of natural language accounts. This richness enables investigators argue about whether or not particular events, such as the omission of TCM-5, are actually causal. There would be no such discussion if everyone accepted the same precise mathematical definition! The key point here is that there must be some form of consistency in determining whether or not to explore particular events during any causal analysis. This can either be done by developing strict mathematical rules that can be applied to formal models of causation. Alternatively, they can be drafted as heuristics that can guide less formal analysis by teams of incident investigators. Different forms of ECF tables might be developed to identify any factors that are particularly important for errors of omission [363]. A further alternative might be to ensure that omitted barriers do not appear in the primary event line of an ECF chart because they are explicitly represented by questions in the causal factor analysis. Unfortunately, the documentation associated with existing applications of the ECF approach does not provide any guidance on how this approach might be developed. Instead, there is an emphasis upon the subjective importance of any analysis. There has been no research to determine whether this results in significant inconsistencies between the analysis of different teams of investigators applying the same technique.

Table 10.12 presents the results from applying ECF analysis to Ground-based software uses imperial not metric units for thruster to compile AMD data file. This event occurred each time an AMD maneuver altered the Climate Orbiter's trajectory. As can be seen, the use of Imperial units stemmed from a failure to follow the Software Interface Specification. This document required the use of metric units but the development staff received insufficient training to appreciate the significance of this document. As with the previous examples of causal factor analysis, this example also shows how the tables can be used to collate information about an event that might otherwise be distributed throughout an ECF chart. In this case, the Software Interface Specification was not used to guide test case generation. This provides an example of the way in which omitted barriers can be represented within the products of a causal factor analysis, rather than being explicitly introduced into an ECF chart as was the case with the decision not to perform TCM-5.

As before, Table 10.12 identifies some of the individuals and groups who were involved in this event. It also refers to a 'mission assurance manager'. This role had existed in previous missions but no-one performed this role during the Climate Orbiter mission. This illustrates how ECF tables can go beyond the omission of barrier events to also represent the lack of key staff who might have prevented the incident. Finally, Table 10.12 identifies some of the features that are shared between a number of similar incidents. In particular, it refers to the role of development documentation in both the Polar Lander and Climate Orbiter case studies. In the former case, requirements document XB0114 failed to provide programmers with enough information about potential failure modes for the Hall Effect sensors. In the later case, software developers failed to follow the Software Interface Specification because they failed to understand the importance wither of this document or the code

Event :	Ground Based Software uses imperial and not metric units for thruster to compile AMD data file
What led to the event?	The project Software Interface Specification was not followed nor was their sufficient oversight to detect the incorrect representation of thruster performance.
What went wrong?	Thruster performance data was encoded in Imperial units in the ground based Small_forces routine. This was used to calculate the values that were stored in the AMD_File. Trajectory modellers within the navigation team used this data. They expected it to be in Metric units. As a result, their calculation of the velocity change from AMD events was out by a factor of 4.45 (1 pound of force = 4.45 Newtons) [570]. Key members of the small forces software team were inexperienced. They needed more training on the ground software development process in general and about the importance of the Software Interface Specification in particular. Inadequate training about end-to-end testing of small forces ground software. Failure to identify that the small forces ground software was potentially 'mission critical'.
How did the barriers fail?	SIS not used to direct testing of the ground software. Unclear if this software underwent independent verification and validation. Management oversight was stretched during transition from development to operations and so insufficient attention was paid to navigation and software validation issues. File format problems with the ground software AMD files prevented engineers from identifying the potential problem. Lack of tracking data.
Who was involved in the event?	Ground software development team, Project management, Mission assurance manager (not appointed).
Is the event linked to a more general deficiency?	Software problems affect Polar Lander. Many of these relate to development documents.

Table 10.12: ECF Analysis of the Climate Orbiter Failure.

that they were writing.

Event	Contextual/ Causal	Justification
Mishap investigation board is established	Contextual	Post-incident event.
Both DS2 probes suffer electrical failure at impact	Causal	The incident would not have happened if this had been avoided.
Forces at impact compromise aft body battery assembly	Causal	The incident would not have happened if this had been avoided. Providing that the RF components were not compromised.
Forces at impact compromise RF components	Causal	The incident would not have happened if this had been avoided. Providing that the battery body assembly was not compromised.
Both DS2 probes impact with the surface	Contextual	Normal or intended behaviour.
Both DS2 probes separate correctly from the MPL	Contextual	Normal or intended behaviour.

Table 10.13: Summary of the ECF Analysis of the Deep Space 2 Incident.

10.3.2 Cause and Contextual Summaries

Causal factor analysis proceeds in the fashion described in previous paragraphs. Investigators iteratively pose counterfactual questions to determine whether each event in an ECF chart can be considered to be causal or not. Table 10.13 summarises the results of this analysis for the loss of the Deep Space 2 probes. As can be seen, there are three causal events: Both DS2 probes suffer electrical failure at impact; Forces at impact compromise aft body battery assembly and Forces at impact compromise RF components. An electrical failure jeopardises the mission if either the aft body battery assembly is compromised or the RF components fail at impact. Each of these events is an element of what Mackie calls a 'causal complex' [508]. It is the conjunction of singular causes within the causal complex that leads to a particular outcome. Crucially, the causal complex is sufficient for the result to occur but it is not necessary. There can be other causal complexes. If any of the necessary causal factors within a causal complex are not present then the incident would not have occurred in the manner described.

Table 10.14 extends the previous analysis of the Deep Space 2 probes to account for the loss of the Polar Lander. This identifies three causal factors. Two are relatively straightforward. This incident would clearly have been avoided if the Hall Effect sensors had not generated transient signals. Similarly, the failure would not have happened if the Lander's engines had not been prematurely cut at 40 meters above the surface. The third event is less easy to assess because it describes the failure of a potential barrier. The software provided some protection against transient signals by rejecting spurious readings from individual sensors. However, it failed to reset the touchdown variable that was used to determine whether the engines should be cut. Table 10.14 argues that this is a causal failure because had the code been written correctly then the incident would not have occurred. This event again illustrates the iterative nature of causal factor analysis.

Even at this advanced stage, it is possible to identify potential improvements to the underlying ECF charts. For example, the analysis presented in Table 10.14 depends on a number of complex

Event	Contextual/ Causal	Justification
Mishap investigation board is established	Contextual	Post-incident event.
Premature Shut-Down of engines (40 meters above surface)	Causal	The incident would not have happened if this had been avoided.
Software marks individual legs as failed if they show spurious signals but does not reset touchdown indicator at 40 meters (entry +5:16)	Causal (Barrier)	The incident would not have happened if this had been avoided. This represents a failed barrier because the software does check for spurious signals in individual legs but does not reset the Touchdown indicator.
Radar detects surface of Mars is 40 meters away (entry +5:15)	Contextual	Normal or intended behaviour.
Software marks a touchdown indicator as true if two spurious signals received from the same leg (10-20 milliseconds after deployment)	Contextual	The incident would not have happened if this had been avoided. The software could have disregarded sensor values until some period after leg deployment.
Transient signals possible from Hall Effect magnets when legs first deploy at 1,500 meters (Entry +4:13)	Causal	The incident would not have happened if this had been avoided.

Table 10.14: Summary of ECF Analysis for Polar Lander Incident (Part 1).

counterfactual arguments. These can be simplified by restructuring the underlying ECF charts. For example, the event labelled Software marks individual legs as failed if they show spurious signals but does not reset touchdown indicator at 40 meters (entry +5:16) can be divided into two component events. One might represent the successful operation of the software defence Software marks individual legs as failed if they show spurious signals. The second event might denote the potential failure Software does not reset touchdown indicator before 40 meters. The former is a contextual event that represents normal or intended behaviour. The latter event can be seen as a causal factor. It represents a failed barrier that might have prevented the incident from occurring had it been correctly implemented.

Table 10.14 summarises the causal and contextual factors that contributed to the loss of the Polar Lander. In particular, it focussed on the potential software failure and its consequent effect of prematurely shutting down the engines while the craft was still some forty meters above the planet surface. Table 10.15 extends this analysis by assessing the events that were used to denote the development and validation of the Lander in previous ECF charts. Two causal events can be identified in this summary: Preliminary design review passed and Launch approved. This analysis again illustrates the practical complexity of counterfactual reasoning about complex failures. For example, it can be argued that both of these events are anticipated within the normal development process and hence should be regarded as contextual rather than causal. The events themselves do not lead to the incident. It is the conjunction of the event together with critical conditions, such as the absence of a system level hazard analysis, that creates a potential failure. Other so-called 'normal' events, such as the end of the cruise phase, are not directly associated with such conditions and hence are not considered to be causal. From this it follows that investigators must not only

Event	Contextual/ Causal	Justification
Last signal from MPL/DS2 (12:02, 3/12/99)	Contextual	Normal or intended behaviour.
Final Trajectory Correction Maneuver (TCM5) begins (05:30, 3/12/99)	Contextual	Normal or intended behaviour.
Cruise phase ends (3/12/99)	Contextual	Normal or intended behaviour.
MPL and DS2 launched (3/1/99)	Contextual	Normal or intended behaviour.
Launch approved	Causal	The incident would not have happened if this had not happened. This could be considered as a normal or intended behaviour. However, the launch should not have been approved without further systems-level analysis and tests.
Development completed	Contextual	Normal or intended behaviour.
Preliminary Design Review passed	Causal	This might be considered a normal or intended behaviour and hence should be contextual rather than causal. However, passing the PDR without further risk management was a causal factor.
Decision to use pulse-mode control	Contextual	This event contributed to the incident because it added to the complexity of the development process and thereby consumed additional design resources.
Decision to use off-the-shelf engines in 4x3 configuration	Contextual	This event contributed to the incident because it added to the complexity of the development process and thereby consumed additional design resources.

Table 10.15: Summary of ECF Analysis for Polar Lander Incident (Part 2).

consider the nature of individual events but also the conditions that affect or modify those events in order to determine whether or not they contributed to the causes of an incident.

Tables 10.16 and 10.17 turn from an analysis of the Polar Lander to examine the ECF charts for the loss of the Climate Orbiter. Table 10.16 identifies a single cause in the events immediately before Mars Orbital Insertion. This relates to the decision not to perform TCM-5. Previous paragraphs have explained how this event can be viewed as causal, if one accepts that TCM-5 is likely to have avoided the incident, or as contextual, if investigators determine that TCM-5 need not have affected the loss of the mission. This illustrates the complexity of informal, subjunctive, counterfactual reasoning. Particular conclusions often depend on the investigators' confidence in a process or device, such as the TCM-5 maneuver. In consequence, the value of structures such as Table 10.16 is not that they simply this difficult form of reasoning. It is, however, that they provide a means of explicitly recording the outcome of such analysis. They also, very importantly, provide a summary justification for any decision to classify an event as either contextual or causal.

Table 10.17 identifies seven causal factors, of which three relate to the failure of potential barriers. The incident would not have occurred if the SM_Forces routines had not used Imperial, rather than Metric, units to calculate the values in the AMD file. These values would not have been so critical if engineers had not rejected to use the barbecue mode or if a symmetrical design had been chosen. The failed barriers relate to the lack of independent verification and validation for the SM_Forces software. They also stem from the limited number of personnel who made the transition between development and operations. The lack of any a priori hazard analysis early in the development project also removed further protection. The identification of these failed barriers as potential causes again depends upon complex forms of counterfactual reasoning. For example, the small number of development staff being moved into operational roles can only be considered a causal factor if investigators believe that a greater number of development staff would have avoided the problems that affected the mission. It is possible to develop formal models that codify and, therefore, simply counterfactual reasoning. However, these approaches ultimately depend upon investigators determining whether or not such changes in the course of events might have avoided the ultimate failure. The complexity of counterfactual reasoning is, therefore, only partly due to the difficulty of constructing valid arguments. It also stems from the inherent difficulty in constructing arguments that are based on limited knowledge about events that we know did not actually take place.

The previous analysis has a number of important limitations. In particular, it follows the recommended ECF practice of focusing the analysis on events [207, 209]. This creates problems because conditions often provide a common link between many different causal events. Such relationships can be represented in an ECF chart. They can, however, become obscured by the tabular form of analysis that is used to summarise the results of any counterfactual analysis. A further concern is that different investigators may make very different choices when deciding whether or not to represent particular factors as events or conditions. For example, we could introduce a condition which states that requirements document XB0114 does not explicitly consider the failure modes for the Hall Effect sensors. The same omission can also be represented by a number of putative events; Requirements document XB0114 published without failure modes or Decision to omit failure modes from XB0114. These concerns are compounded by the observation that managerial failures are often represented as conditions while individual instances of human error often reveal themselves as discrete events.

A number of approaches can be used to counter-balance this bias towards events. For instance, it is possible to repeat the previous analysis but instead focus upon conditions rather than events. An example of the counterfactual question would then be 'would the incident have occurred if it was not the case that the Climate Orbiter's ground software development staff had limited training in this application domain?'. This approach offers a number of benefits. In particular, it ensures that investigators revisit the many different conditions that can emerge during the previous stages of analysis. This process of cross-checking can help to reveal instances in which the same conditions effect many different aspects of an incident. This approach can, however, also introduce a number of practical difficulties. Almost all of the counterfactual questions that can be applied to the conditions in an ECF chart follow the subjunctive forms that have frustrated our previous analysis of failed barriers. It is very difficult to derive an objective answer to the previous example. How can we determine whether improved training would have avoided the incident? An alternative approach is

Event	Contextual/ Causal	Justification
MCO Mishap Investigation Board is formed (15/10/99)	Contextual	Post-incident event.
Operations navigation team consult with spacecraft engineers to discuss discrepancies in velocity change model (27/9/99)	Contextual	Post-incident event.
Last signal from MCO (09:04:52, 23/9/99)	Contextual	Normal or intended behaviour. The signal was lost as the craft passed behind the planet during orbital insertion.
Mars Orbital Insertion begins (09:00:46, 23/9/99)	Contextual	Normal or intended behaviour.
Cruise phase ends (23/9/99)	Contextual	Normal or intended behaviour.
TCM-5 is discussed but not executed (16-23/9/99)	Causal (Barrier)	The failure of a barrier causes problems for counterfactual reasoning because it relies upon subjunctive arguments that may, or may not be justified. In this case, we consider it likely that TCM-5 would have avoided the incident had it been performed.
(File format) anomaly is not reported through Incidents, Surprises, Anomaly system	Contextual (Barrier)	This also depends on a subjunctive argument about whether or not the ISA system might have prevented the incident had it been used. In this case, it is considered that the incident might still have occurred even if the file format anomaly had been reported.
It is apparent that AMD file data is anomalous (N + 7/4/99)	Contextual	Not causal because it created an opportunity to avoid the incident.
File format problems for AMD data is corrected (N/4/99)	Contextual	Not causal because it created an opportunity to avoid the incident.

Table 10.16: Summary of the ECF Analysis of the Climate Orbiter Incident.

Event	Contextual/ Causal	Justification
Ground-based software uses Imperial and not metric units for thruster to compile AMD data file	Causal	The incident would not have happened if this had been avoided.
Limited independent testing of the ground-based SM_Forces routines	Causal (Barrier)	It is considered likely that the incident would not have occurred if there had been greater independent testing of these routines.
SM_Forces routines are written using imperial and not metric units for thruster performance	Causal	The incident would not have happened if this had been avoided.
Angular Momentum Desaturation events	Contextual	Normal or intended behaviour given the MCO's asymmetric design and the decision to reject the barbecue maneuver.
Systems engineering decision to reject daily 180 degree flip to cancel angular momentum build-up.	Causal	The incident might not have happened if the engineers had decided to perform the 'barbecue' maneuver. However, there remains a degree of doubt that this further navigation problems might have been introduced or gone undetected.
Systems engineering decision to use a solar array that is asymmetrical to the MCO body	Causal	The incident might not have happened if a symmetrical design had been introduced similar to the Global Surveyor.
MCO launch (11/12/98)	Contextual	Normal or intended behaviour.
Minimal number of development staff transition to operations (11-12/98)	Causal (Barrier)	The incident might not have happened if more staff had moved from development to operations.
Decision not to perform an a priori analysis of what could go wrong on the MCO.	Causal (Barrier)	The incident might not have happened if more thought had been given to the problems involved in using the MCO design to achieve the navigation accuracy required by the mission.

Table 10.17: Summary ECF Analysis for Climate Orbiter Incident (Part 2).

to use Causal-Context summaries as a form of index into the underlying ECF charts. These diagrams retain the broader conditions that help to shape the context for any incident. In contrast, the summary tables strip out much this detail to focus on the elements of Mackie's causal complexes. Cause-context summary tables and ECF charts together provide a stepping stone towards any subsequent root cause analysis. The following paragraphs address a number of the key issues that must be addressed by any root cause analysis technique.

When to begin? Previous chapters have also argued that the early stages of an investigation are often guided by investigators' working hypotheses about the causes of an incident. It is important, however, that these informal ideas should be explicitly represented relatively early if finite investigatory resources are to be maximised. This requirement must be balanced against the dangers of biasing an investigation towards certain causes. Root cause analysis uses the results of the previous techniques to identify common factors behind causal events. As noted in the previous paragraphs, these common factors may already have been identified as conditions within an ECF chart. It is important to stress, however, that root cause analysis "is not an exact science" [207]. The processes of analysis and investigation often uncover potential root causes that were not considered during previous stages of analysis. It is important, therefore, not to freeze the ECF chart or the cause-context tables during the early stages of any analysis.

How do we validate the analysis? We have argued that ECF charts and cause-context diagrams are 'living' documents that must be updated as new information becomes available. It is important, however, that investigators validate the products of any causal analysis. Typically, this is done through regular, minuted team meetings. Increasingly these are used to approve the publication of draft analysis documents via organisational intranets. They provide shared resources that help to guide the continuing investigation. Such publication and distribution mechanisms help to coordinate investigators' activities but must be protected from public disclosure. Ultimately, the products of any root cause analysis must be approved by the members of an investigation team before a final report can be written. This mechanisms for achieving this agreement depend on the scale of the incident reporting system. In local applications, there may only be a single individual who is available to perform the analysis and draft the report. In larger systems, however, there may be formalised procedures for 'signing off' the products of any root cause analysis. These procedures can involve higher levels of management. This raises serious practical and ethical issues if this final stage of approval is seen as a means of potentially filtering the results of any analysis. Some organisations have guarded against this by allowing senior management only to annotate root cause analyses. They are prevented from altering what has already been written. While this approach offers some protection against undue influence, it does not guard against the myriad of informal pressures that can be brought to bare on an investigation team.

How many root causes? The Department of Energy guidelines state that investigators should identify at least one but probably not more than three or four root causes [207]. This guideline seems to be derived from the pragmatics of incident investigation within particular industries. They do not, however, provide any justification for their suggestion. This is unfortunate. Such a pragmatic limit can be seen as a barrier to organisational learning from any mishap in which there were more that four root causes. Such concerns are exacerbated by the observation that there are often many different ways for an incident to occur. In consequence, there any incident investigation may yield a number of root causes for each of these different scenarios. For instance, the Polar Lander could have been lost because of the premature shut-down of the engines. It might also have been caused by a failure in the separation of the Deep Space 2 probes and the Lander from the cruise stage. It could have been caused by a landing on unfavourable terrain. It might also have been caused by failure in the communications up-link and so on. Each of these scenarios was considered to be plausible by the NASA investigation team. Although each hypotheses yielded a small number of root causes, the cumulative effect of considering many different failure scenarios helped the investigators to identify a significant number of lessons for future missions. This would not have been possible had they stopped at the four or five root causes recommended above. It seems more profitable to view resource constraints as the limiting factor. The extent of any root cause analysis provides a good indication of the perceived criticality of any potential failure.

What are the parameters of the analysis? The ECF guidelines argue that "the intent of the analysis is to identify and address only the root causes that can be controlled within the system being investigated, excluding events and conditions that cannot be reasonable anticipated and controlled, such as natural disasters" [207]. It is clearly difficult to control natural disasters, however, this wide ranging approach does pose a number of important questions. Previous sections have explained how many local incident reporting systems 'target the doable'. This can prevent effective action from being taken to address common problems that might affect a number of different local groups. In particular, managerial and organisation constraints may be viewed as outside the control of operational departments. It is, therefore, important that any root cause analysis technique should provide explicit means of addressing these higher-level causes of failure.

The previous paragraphs have described some general attributes of the root cause analysis. They have not, however, provided any guidance about the methods and techniques that might be applied to identify these factors from the mass of information that can be derived from the previous stages of analysis. The following sections, therefore, present two different techniques that can be used to identify root causes from the events and conditions that are described in ECF charts and cause-context tables.

10.3.3 Tier Analysis

Tier diagramming is a root cause analysis technique that focuses on those levels of management that have the responsibility to correct potential problems. It is one of several techniques, including Pate-Cornell's 'Accident Analysis Framework', that exmplicitly force investigators to consider organisational factors as the initial root causes of many failures [665]. Each row in one of these diagrams refers to a different level of management within an organisation. They are intended to represent levels of organisational responsibility that range from the operator up to senior management. The columns in a tier diagram list the causal factors that are derived from the Causal factor analysis together with any higher-level root causes that may or may not be identified. This is illustrated by Table 10.18. It is important to note, however, that this a generic template that must be tailored to reflect the organisations that are involved in a particular incident. Each causal factor is assigned to a tier of management responsibility. This is intended to help identify any common links between causal factors that relate to particular levels in an organisation. For instance, a failure in supervision would be exposed by a number of causal factors that cluster around this level in the tier diagram. This is intended to offer a number of benefits to any incident investigation. In particular, it helps to focus any root cause analysis on the deeper organisational causes of failure [701]. The tabular format also helps to structure an investigation around concepts, or groups, that have a clear organisational meaning for those involved in an incident. This is important because many incident reports often talk in vague terms about a 'failure in safety culture' without grounding these observations in the activities of particular organisations and groups. A further benefit is that responsibility is explicitly assigned for each root cause and causal factor. These judgements provide a focus for subsequent discussion and can, ultimately, help to form the recommendations for future practice.

Tier	Causal Factors	Root Cause
5: Senior Management		
4: Middle Management		
3: Lower Management		
2: Supervision		
1: Workers Actions		
0: Direct Cause		

Table 10.18: Format for a Tier Diagram [207].

Different tier diagrams are drawn up for each of the organisations that is involved in an incident. In our case studies, therefore, we would anticipate separate tier diagrams for NASA Headquarters and for NASA JPL and for the subcontractor LMA. It is also possible to refine such diagrams to look at different groups and teams within each organisation. For instance, it is possible to distinguish management tiers within the development process of the Climate Orbiter from operation groups. Tier diagramming, typically, begins with the organisation that is most closely involved in the incident. The first diagram in both the Polar lander and Climate Orbiter case studies would focus on the LMA operational teams. Further diagrams would then represent the contractor organisation for which LMA was subcontracting. In particular, tier diagrams should also represent any organisations that are involved in the oversight or regulation of the contractor's and subcontractor's activities. Tier diagramming, therefore, has two prerequisites. Firstly, investigators must have already identified a number of potential causal factors using techniques such as causal factor analysis. Secondly, they must also have a clear understanding of the management structures that characterise the organisations involved in an incident. Once this information is available, the analysis proceeds in the following stages:

1. Develop the tier diagram. Create a tier diagram that reflects the management structure of the organisation being considered.

2. Identify direct causes. Examine the cause-context summaries to identify any catalytic events that cannot be directly associated with operators or management activities. Enter these along the direct cause row, shown in Table 10.18. Repeat this process for any conditions that are associated with these causal events in an ECF chart. Initially, this tier might contain events that describe the failure of process components or problems due to the contamination of raw materials. As analysis progresses, however, it is likely that most of these direct causes will be associated with other tiers in the diagram. For instance, component failures may be due to a managerial failure to ensure an adequate maintenance regime. Similarly, the contamination of raw materials can be associated with acquisitions and screening policies.

3. Identify worker actions. For each causal factor in the cause-context summary, ask whether or not they stemmed directly from 'worker actions'. A number of guidelines can be proposed to direct this stage of the analysis. For instance, the US Department of Energy has developed a number of questions that are intended to help determine whether or not a causal factor should be associated with worker actions [207]. These include whether or not the worker's knowledge, skills and abilities were adequate to perform the job safely. They also ask whether the worker understood the work that was to be performed. As with direct causes, these actions often raise questions about the performance of other groups in a tier diagram. The worker's lack of understanding may be due to an inadequate training regime. Investigators must, therefore, ask whether or not the worker was solely responsible for the causal factor. If the answer is no then investigators must move the event to a higher tier in the diagram. As before, investigators must also introduce any associated conditions into a tier diagram if they provide necessary additional information about causal events.

4. Analyse remaining tiers. The analysis progresses in a similar fashion for each tier. The intention is to place each causal factor as high up the diagram as possible. Ultimately, as we have seen, all incidents can be associated with regulatory problems or a failure in oversight. It is important, however, to balance this observation about ultimate responsibility against the need to identify those levels in an organisation that are most directly responsible for certain causal factors. As mentioned in the previous paragraph, this is most often done by developing analytical guidelines. These guidelines help investigators to assess whether or not a causal factor can be associated with a particular tier in the diagram. They are, in turn, typically derived from the safety cases that justify the operation of an application process. For instance, if middle management has an identified responsibility to ensure the operation of an incident reporting system then it is possible to place any causal factor that relates to the failure of such a system at this level in a tier diagram.

5. Identify links. After all of the causal factors and associated conditions have been entered into a tier diagram, investigators can begin to look for common factors. As with the previous stages in this form of analysis, the success of this activity depends upon the skill and expertise of the

investigator. This, in turn, can have a profound impact on the course of any investigation. As Lekberg notes, the previous background and training of an investigator can have a profound impact on the results of their analysis [484]. The key point is not, however, to eliminate these individual differences but to use the tier diagram as a means of explicitly representing the key stages in any root cause analysis. Other investigators can then inspect these diagrams to identify other connections between causal factors or, if necessary, to argue against proposed links. Investigators can use different colours or symbols to denote those causes that are considered to be linked.

6. Identify root causes. Compare each of the causal factors in the tier diagrams against the definition of a root cause. A root cause is distinguished by Lewis' counterfactual argument that if A and B are states (conditions) or events, then A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either [490]. This is essentially the same requirement that was used to distinguish causal from contextual factors in the causal factor analysis. They can also be thought of as causal factors that, if corrected, would prevent recurrence of the same or similar incidents. We would also impose an additional requirement based on Mackie's distinction between general and singular causes [508]. Root causes must address a class of deficiencies, rather than single problems or faults. Correcting a root cause not only prevents the same incident from recurring but also solves deeper line management, oversight and management system deficiencies that could cause or contribute to future mishaps [207]. If a causal factor meets these criteria then an additional entry can be made to denote this finding in the third table of the tier diagram, illustrated in Table 10.18. Investigators must, therefore, compose a root cause 'statement' to summarise each of the causal factors groupings that were identified in the previous stage of analysis.

Root cause analysis can reveal events and conditions that were not represented on ECF charts, ECF tables or cause-context summaries. . These must be added to ensure consistency between these various products of a root cause analysis. It should also be noted that one tier diagram may provide input for another. For instance, if the upper management of a contractor was responsible for a particular root cause then the regulator and supervisory organisation may share responsibility for that particular root cause if there is a deficiency in the directives given by those organisations.

The remainder of this section applied the tier diagramming approach to identify root causes for both the Polar Lander and the Climate Orbiter case studies. This analysis begins by identifying the relevant management and organisation structures that were involved in this incident. The Mars Independent Assessment Team have provides information about the internal management structures within NASA headquarters and within JPL [569]. Unfortunately, it can be less easy for investigators to obtain detailed information about subcontractors' management structures even in the aftermath of a serious incident. The subsequent analysis, therefore, must also exploit a number of inferences about the reporting structures that characterised the day to day operation of the Mars Surveyor projects.

Figure 10.26 illustrates the complexity of the management structures that were involved in the Mars Program at NASA Headquarters. Not only do such organisational features complicate any tier analysis, they also had a significant impact on the loss of the Polar Lander and the Climate Orbiter. During the initial formation of the program, the JPL Program Manager had to deal with the Advanced Technology and Mission Studies Division. During implementation, they interacted with the Mission and Payloads Development Division. For the operational phase of the program, the JPL Program Manager dealt with the Research and Program Management Division. During all of this the manager must also interact with the Science Board of Directors. These various channels of communication between NASA headquarters staff and the JPL Mars Program Manager caused problems for both organisations. The independent assessment team found that "ineffective communication between JPL management and NASA Headquarters contributed to an unhealthy interface and significant misunderstandings in conducting the Mars Surveyor Program" [569]. NASA Headquarters believed that they were articulating program objectives, mission requirements and constraints. JPL management interpreted these as non-negotiable demands over costs, schedules and performance requirements. Concern about losing contracts and funding also prevented JPL

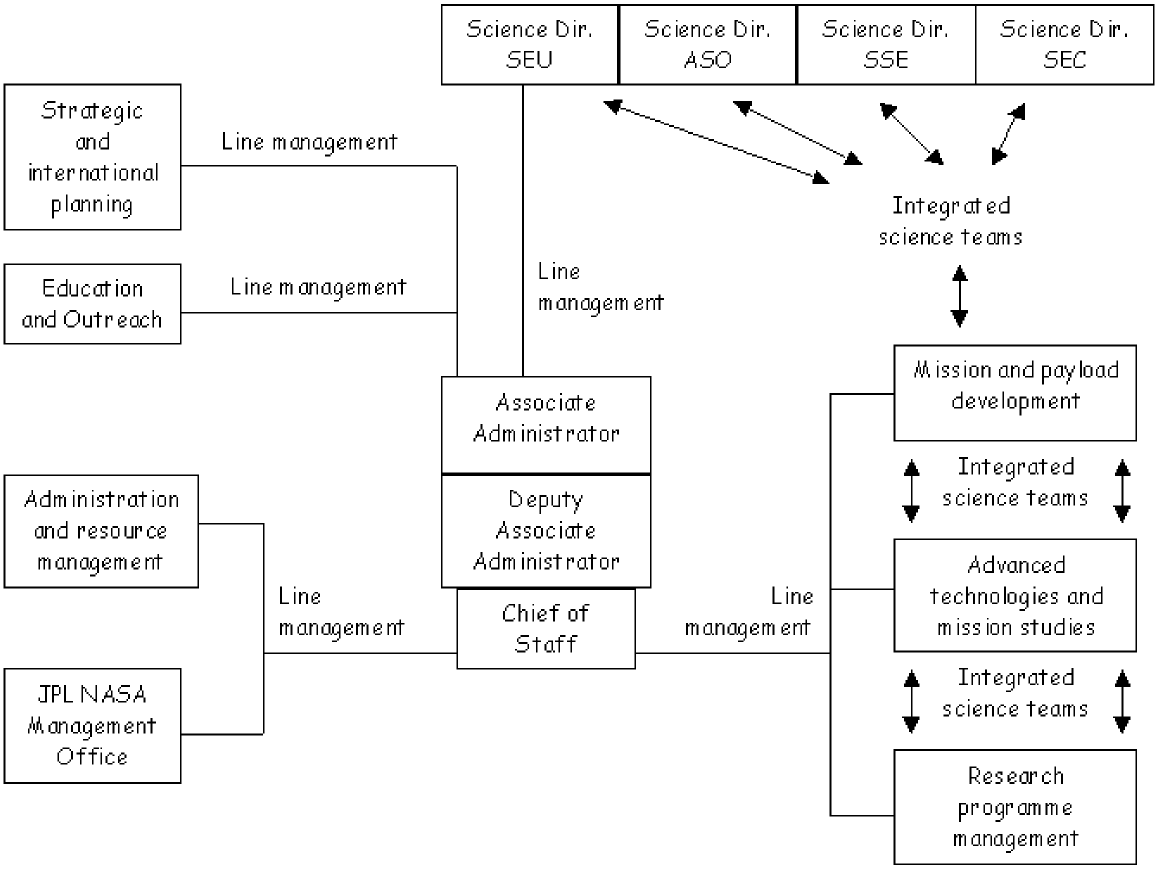

Figure 10.26: NASA Headquarters' Office of Space Science [569]

management from effectively express their concerns to NASA Headquarters about programmatic constraints. The independent assessment team also concluded that NASA Headquarters did not seem receptive to receiving bad news.

JPL's Mars Program Office initiated the Mars 98 project and was responsible for planning, program advocacy and flight project development between 1994 and 1996. The roles and responsibilities of this office were, however, interpreted differently in the JPL Mars Program Office and the NASA Headquarters sponsoring office. This led to several conflicts about mission direction that ultimately diverted management resources away from mission development. These difficulties illustrate an important practical barrier to tier analysis. One of the precursors to an incident may be the breakdown of management structures. The roles and responsibilities of each level of the table may, therefore, be very difficult to distinguish: "individual projects were not developed or managed within a clearly defined overall framework that identified interdependencies and risk management strategies" [569].

In 1996, NASA Headquarters delegated full program management authority to the NASA Centers. JPL, therefore, created a Mars Exploration Directorate that reported directly to the Laboratory Director. This directorate assumed responsibility for program management and assumed most of the duties that have previously been associated with the NASA Headquarters sponsoring office. One consequence of this reorganisation was that JPL's Mars Exploration Directorate lost a single point of contact at Headquarters. In August 1996, the management structure of the Mars programs was further complicated by the announcement that potential signs of life had been found on a meteorite that was assumed to have come from Mars. The heightened public interest led to further changes in JPL's organisation. An increased emphasis was placed on robotic exploration to support the long-term needs of Human Exploration. These missions were managed by a different part of

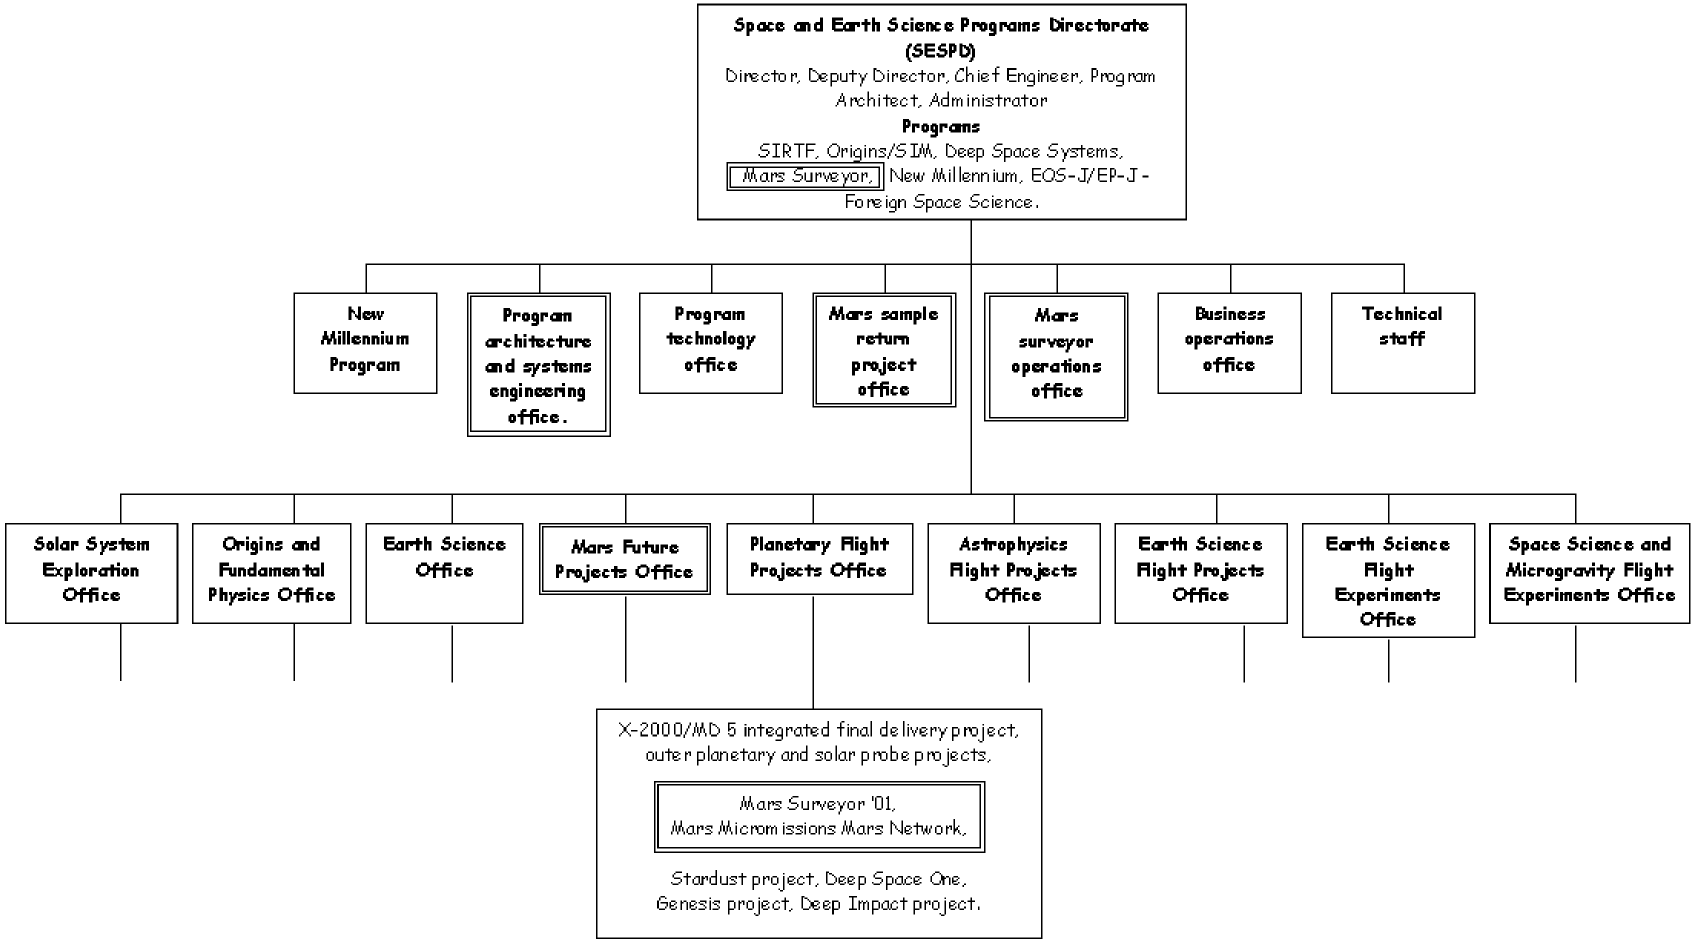
Space and Earth Science Programs Directorate
(SESPD)
Director, Deputy Director, Chief Engineer, Program Architect, Administrator
Programs
SIRTF, Origins/SIM, Deep Space Systems,
Mars Surveyor, New Millennium, EOS-J/EP-J-Foreign Space Science.
New Millennium Program
Program architecture and systems engineering office.
Program technology office
Mars sample return project office
Mars surveyor operations office
Business operations office
Technical staff
Solar System Exploration Office
Origins and Fundamental Physics Office
Earth Science Office
Mars Future Projects Office
Planetary Flight Projects Office
Astrophysics Flight Projects Office
Earth Science Flight Projects Office
Earth Science Flight Experiments Office
Space Science and Microgravity Flight Experiments Office
X-2000/MD 5 integrated final delivery project, outer planetary and solar probe projects,
Mars Surveyor '01, Mars Micromissions Mars Network,
Stardust project, Deep Space One, Genesis project, Deep Impact project.

Figure 10.27: JPL Space and Earth Sciences Programmes Directorate [569]

Headquarters

JPL responded to these changes in priorities by partially reorganising its own management structure in 1998. This was followed by wider changes in 1999. JPL amalgamated its space and Earth science teams into a single directorate. The intention was to coordinate the management of an increased number of programs and projects in both of these areas. The Mars Program Manager no longer reported to the Laboratory Director as a separate, independent entity. Project managers were to report at a lower level. Figure 10.27 illustrates the organisational structure of the JPL Space and Earth Sciences Programs Directorate after the 1999 reorganisation. The Mars projects are shown among sixty-eight other projects in the third tier of management. They are, therefore, isolated from the direct reporting structures of senior JPL management. Although Figure 10.27 represents the 1999 reorganisation, the independent assessment team argued that this reflects the project isolation that contributed to the failure of the Mars'98 project.

The previous paragraphs have summarised the management structures within NASA headquarters and within JPL. They have also argued that the dynamism of many organisations can create significant problems when applying tier analysis to real-world management structures. The different teams and individuals who are associated with different levels in a tier diagram may change as organisations attempt to adapt to the pressures that are created by many high-technology projects. One solution would be to develop a number of tier diagrams to represent these different changes in project management. An alternative approach is to exploit a relative abstract classification of organisational structures, similar to those shown in Figure 10.18 and then provide more detailed information to support the interpretation of those categories at particular stages of the incident.

A number of further challenges complicate the development of tier diagrams. In particular, it may not be possible for the investigators from one organisation to gain access to detailed information about the management of another organisation. As we have seen, it is relatively easy to access documentation about NASA management structures. It is far harder to find comparable information about the organisation of the commercial subcontractors. In consequence, investigators may be forced to exploit the more generic tiers that were introduced in Table 10.18. Even if this approach is exploited, investigators face a number of further problems. For example, if there are several organisations involved in an incident then they must determine which causes relate to which tier diagram. This can partly be based on any existing project documentation, however, it also requires considerable skill and judgement on the part of individual investigators. For example, the following quote illustrates how LMA were responsible for the development of the Mars Surveyor programme. JPL staff were involved in some of these activities but they also provided higher level management functions:

> "The Mars Surveyor program'98 Development Project used a prime contract vehicle to support project implementation. Lockheed Martin Astronautics (LMA) of Denver, Colorado was selected as the prime contractor. LMA's contracted development responsibilities were to design and develop both spacecraft, lead flight system integration and test, and support launch operations. JPL retained responsibilities for overall project management, spacecraft and instrument development management, project system engineering, mission design, navigation design, mission operation system development, ground data system development, and mission assurance. The MSP 98 project assigned the responsibility for mission operations systems/ground data systems development to the Mars Surveyor Operations Project, LMA provided support to Mars Surveyor Operations Project for mission operations systems/ground data systems development tasks related to spacecraft test and operations." [564]

This quotation illustrates the practical difficulties that are involved in separating out the responsibility that each organisation might assume for certain causes of safety-critical incidents. In consequence, the following tables represent one particular viewpoint. They act as a focus for subsequent discussion rather than a unique assignment of causal factors to particular management layers in each of the organisations.

Figure 10.19 provides an initial assignment of causes to various layers within the contractor organisation. In addition to these causal factors, identified in the cause-context summaries, it is also possible to introduce conditions that are also perceived to have contributed to the incident. As mentioned, these conditions can represent longer term factors that cannot easily be represented as discrete events and so may be overlooked by the previous forms of analysis. For instance, previous ECF charts identified the way in which some project requirements were not passed on in sufficient detail. This was shown as a condition labelled Requirements are not passed on in sufficient detail nor are they backed by an adequate validation plan in Figure 10.11. This created problems because individual project managers had to interpret what was admissible in pursuit of the objectives set by Faster, Better, Cheaper. Figure 10.19, therefore, introduces a number of similar conditions into the tier diagram.

It is important to note that Figure 10.19 represents the management structure that was in place at JPL between 1994-1996. It was during this period that JPL's Mars Program Office initiated the Mars 98 project and was responsible for planning, program advocacy and flight project development. As noted in previous sections, tier analysis is complicated by the fact that the management tiers were altered several times during the project lifecycle. Figure 10.27, shown previously, illustrates the JPL management structure that was put in place from 1996. A new Mars Exploration Directorate was created within JPL to coordinate many of the activities that were previously performed by NASA Headquarters and so are not considered in Figure 10.19.

Figure 10.19 illustrates the way in which tier analysis tends to associate root causes with the higher levels of management. This is a natural consequence of the iterative process that is used to analyse each causal factor; the intention is to place each causal factor as high up the diagram as possible. This is an important strength of the technique. The investigators' attention is drawn away from individual instances of operator error. Undue emphasis may, however, be placed on individuals at higher levels within an organisation. This is inappropriate if operational responsibility is devolved to lower levels within the management structure. Under such circumstances, any root cause for the failure might have to be associated with several different levels within an organisation.

The distribution of responsibility within an organisation is illustrated in Figure 10.19 by root causes at both senior and middle management level. Although senior personnel provided insufficient guidance on the implementation of NASA's Faster, Better, Cheaper strategy, middle management might still have fought to obtain adequate resources. This also illustrates the subjective nature of tier analysis. It can be argued that these two root causes are so closely linked that they should be amalgamated into a single higher-level description. If Senior Management had provided strong guidance about the implications of the Faster, Better, Cheaper strategy for design and validation then Middle Level Management would have had less need to fight for additional resources. On the other hand, it can be argued that these root causes should be distinct because Senior Management must rely on their colleagues to provide adequate information about the operational implication of accepting such tight resource constraints. Similarly, there are some causal factors in Figure 10.19 that could have been represented as root causes. The decision not to implement TCM-5 is an example of one such event. If this maneuver had been implemented then the incident could have been avoided. The lack of preparation for this maneuver and the consequent decision not to implement it might, in combination with other factors, lead to future incidents. The key point here is that either approach would represent a valid application of tier analysis. The output of this process depends upon the skill, expertise and viewpoint of the investigator. It, therefore, must be carefully validated by peer review. One means of validating our analysis would be to compare Figure 10.19 with the output of an independent tier analysis performed by another investigator. There may, however, be more general biases that are introduced by the use of this particular form of analysis. An alternative means of validating these findings is to compare the results of our analysis with those obtained by investigators using other approaches. For example, the following section will repeat the analysis of our case studies using Non-compliance classifications. For now it is sufficient to summarise the findings of the Mars Program Independent Assessment Team Report. They used a range of less structured techniques to derive the following conclusions about contractor involvement in the root causes of the incident:

> "(NASA, JPL, and LMA) have not documented the policies and procedures that make up their Faster, Better, Cheaper approach; therefore, the process is not repeatable. Rather, project managers have their own and sometimes different interpretations. This

Tier	Causal Factors	Root Cause
Senior Management	Requirements are not passed on insufficient detail nor are they backed by an adequate validation plan. Decision not to perform an a priori analysis of what could go wrong on the MCO. Limited independent testing of the ground-based SM_Forces routines.	No documented guidance on implementing Faster, Better, Cheaper prevented project managers from resisting pressures to cut costs/schedules that might compromise mission success.
Middle Management	Minimal number of development staff transition to operations (11-12/98). SM_Forces routines are written using imperial and not metric units for thruster performance.	Lack of resources for the Mars Surveyor Program limited the number of staff available and may also have prevented those staff from receiving adequate training on critical aspects of the mission.
Lower Management	TCM-5 is discussed but not executed (16-23/9/99)	
Supervision		
Workers Actions	Systems engineering decision to reject daily 180 degree flip to cancel angular momentum build-up. Systems engineering decision to use a solar array that is asymmetrical to the MCO body	
Direct Cause	Ground-based software uses Imperial and not metric units for thruster to compile AMD data file	

Table 10.19: LMA Tier Diagram for the Climate Orbiter Mission.

> can result in missing important steps and keeping lessons learned from others who could benefit from them... Mars 98 had inadequate resources to accomplish the requirements. Through a combination of perceived NASA Headquarters mandates and concern for loss of business, JPL and LMA committed to overly challenging programmatic goals. The JPL management perception was that no cost increase was permissible and the aggressive pricing strategy adopted by LMA exacerbated the problem. The pressure of meeting the cost and schedule goals resulted in an environment of increasing risk in which too many corners were cut in applying proven engineering practices and the checks and balances required for mission success... Inadequate project staffing and application of institutional capability by JPL contributed to reduced mission assurance. Pressure from an already aggressive schedule was increased by LMA not meeting staffing objectives early in the project. This schedule pressure led to inadequate analysis and testing. An additional important role for senior management, whether at NASA, JPL, or LMA, is to ensure the establishment of, and compliance with, policies that will assure mission success. For example, these policies should address design (at the component, system, and mission life cycle level), test and verification, operations, risk management, and independent reviews." [569]

As can be seen, several of the themes identified by the Mars Program Independent Assessment Team mare summarised as root causes in the tier analysis of Figure 10.19. There are some differences. In particular, the team's report brings together many of the factors that we have identified and links them to the contact management's perception of project risk. Our analysis was performed prior to reading this document. With this additional insight, however, it would be possible to reformulate the previous diagram to reflect these more general concerns. This again reflects the point that root cause analysis is an iterative process. ECF charts, cause-context summaries, tier analysis are all artifacts that help to document the path towards a causal analysis. They do not replace the skill and expertise of the investigators nor do they 'automate' key stages of the analysis.

Figure 10.20 builds on the previous analysis by examining the root causes of the Climate Orbiter failure from the perspective of the JPL management structure. Unlike the contractor organisations, more can be identified from the published documentation about management structures within this organisation. As mentioned previously, JPL retained responsibilities for "overall project management, for spacecraft and instrument development management, for project system engineering, mission design, navigation design, mission operation system development, ground data system development and mission assurance" [570]. From this is follows that JPL staff were ultimately responsible for the development and testing of the navigation software. It can, therefore, be argued that some of the causal factors associated with navigation systems development should be removed from Figure 10.19 The contractor was not responsible for overseeing this aspect of the mission. These factors have been retained because the NASA investigators commented on the difficulty of making such precise distinctions, staff often could not reply to questions such as 'who is in charge?' or 'who is the mission manager?' [570].

Figure 10.20 shows how causal factors affect several of the organisations that are involved in any incident. This diagram presents many of the events and conditions that were identified in the tier analysis for LMA staff. However, the supervisory and managerial role of JPL staff is reflected by the way in which many of these causal factors are associated with different levels in the management structure. For instance, the event TCM-5 is discussed by not executed was associated with lower levels of management within the contractor organisation but is associated with the program management in JPL. The Flight Operations Manager should have polled each subsystem lead to ensure that they had reviewed the data and believed that the Climate Orbiter was in the proper configuration for the event. [570] However, this protocol had not been developed nor had any manager been explicitly identified to lead this decision making process. It might, therefore, be argued that responsibility rested with the JPL program manager, as shown in Figure 10.20.

Figure 10.20 also illustrates the manner in which tier analysis can expose different root causes for similar causal factors within different organisations. For example, the inadequate risk analysis and the lack of development staff who transitioned into operations might indicate a degree of complacency on the part of the JPL management team. The NASA investigators found evidence of a perception at

Tier	Causal Factors	Root Cause
5: Senior Management (JPL Laboratory Director and Mars Program Office Director)	Minimal number of development staff transition to operations (11-12/98) Limited independent testing of the ground-based SM_Forces routines Decision not to perform an a priori analysis of what could go wrong on the MCO.	Feeling that orbiting Mars in routine. Insular relationship with LMA prevented adequate risk assessment and mitigated against independent reviews.
4: Middle Management (Climate Orbiter Project Manager)	TCM-5 is discussed but not executed (16-23/9/99)	
3: Lower Management (Flight Operations Manager/Flight Development Manager)	SM_Forces routines are written using imperial and not metric units for thruster performance. Systems engineering decision to reject daily 180 degree flip to cancel angular momentum build-up. Systems engineering decision to use a solar array that is asymmetrical to the MCO body	

Table 10.20: JPL Tier Diagram for the Climate Orbiter Mission.

JPL that "orbiting Mars is routine" [570]. This perception was based on previous mission successes. However, it resulted in inadequate attention being paid to navigation risk mitigation.

Figure 10.20 also illustrates the way in which tier diagram must account for the relationship between the management structure that is being considered and any other organisations that are involved in an incident. In this case, the insular relationship between JPL and the contract organisation is identified as a root cause behind the lack of independent testing and inadequate risk assessment. This analysis raises a number of structural properties about our use of the tier diagrams in Figure 10.20. As can be seen, causal factors and root causes are associated with different levels of management. No distinction is made between these causes. For instance, only two out of the three causal factors at the top levels of the JPL management structure are associated with the insularity, mentioned above. Similarly, we have not shown how causal factors at various levels in a tier diagram might contribute to a root cause. Additional annotations could be introduced to represent this information. Care must be taken if the resulting diagrams are not to become illegible.

As before, we can compare the results of the tier analysis with the findings of the Mars Program Independent Assessment Team. The root cause analysis illustrated in Figure 10.20 is based on a subset of the evidence that was available to this investigation team. Our analysis was, however, done prior to reading their account:

> "The JPL/Lockheed Martin Astronautics interface for Mars 98 was characterised by a positive, close working relationship between the JPL and LMA project managers and their offices. However, this relationship had a negative, insular effect when accepting excessive risk... Inadequate project staffing and application of institutional capability by JPL contributed to reduced mission assurance. Pressure from an already aggressive schedule was increased by LMA not meeting staffing objectives early in the project. This schedule pressure led to inadequate analysis and testing... The team found multiple examples of ineffective risk identification and communication by both JPL and LMA. Compounding this, JPL and LMA each deviated from accepted and well-established engineering and management practices. Risk identification and any significant deviations from acceptable practices must be communicated to the customer in an open, timely, and formal fashion." [569]

It is difficult in the aftermath of such an incident to be sure that this analysis has not biased my interpretation of the incident. The findings of the Mars Program Independent Assessment Team were publicised in press accounts. They are also referenced in the pages that provided access to on-line versions of primary sources that were used in our analysis. Any comparison between the results of our tier analysis and the assessment team's report cannot, therefore, be regarded as an independent or formal validation of the root causes analysis. In contrast, Figure 10.20 simply illustrates that it is possible for some of the independent assessment team's findings to be represented within a tier diagram. It is also important to identify the differences between our ECF/tier analysis and the findings of the independent assessment team. In particular, the root causes in Figure 10.20 do not address the communications problems that existed between JPL and NASA headquarters. The Mars Program's Independent Assessment Team report emphasised that these problems prevented JPL management from gaining a clear understanding of the resource implications behind the Faster, Better, Cheaper strategy. These concerns are, however, represented in Table 10.21 that presents a tier analysis of NASA headquarter's involvement in the loss of the Climate Orbiter.

Figure 10.21 illustrates the way in which investigators can use both the conditions and the events in an ECF chart to support any subsequent tier analysis. In this case, NASA headquarters had little direct involvement in the events that led to the loss of the Climate Orbiter. Investigators would, therefore, have considerable difficulties in constructing a root cause analysis that was based solely upon such direct involvement. In contrast, it can be argued that NASA headquarters played an important role in establishing the conditions that led to this incident. Figure 10.21 therefore goes beyond the causal events that were considered in previous tier diagrams to look at the conditions that were identified in early ECF charts of the Climate Orbiter incident, such as Figure 10.8. This example is typical of tier diagrams that consider the role of regulatory or supervisory organisations in such failures. It is also important to note that such factors are often omitted from some reports of

Tier	Causal Factors	Root Cause
5: Senior Management (Board of Directors, Science)	Project oversight problems stem from complex relationship between JPL and LMA (and NASA HQ)	Failure to communicate the mission implications of the Faster, Better, Cheaper strategy.
4c: Middle Management (Associate Administrator, Office of Space Science)		
4b: Middle Management (Science Chief of Staff)	Lack of managerial leadership in promoting responsible attitudes to Incidents, Surprises and Anomaly reporting	Failure to communicate the importance of expressing concerns both about specific implementation issues as well as resource/management problems.
4a: Middle Management (Advanced Studies Division, Mission Development Division, Research and Program Management Division etc)	Requirements are not passed on in sufficient detail nor are they backed by an adequate validation plan	

Table 10.21: NASA HQ Tier Diagram for the Climate Orbiter Mission.

an incident. For example, the initial report into the Climate Orbiter contained no reference to the involvement of NASA headquarters at all [564]. This is justified by the initial focus on the direct causes of the incident. The subsequent report into Project Management in NASA by the Mars Climate Orbiter, Mishap Investigation Board only contained four references to NASA headquarters [570]. None of these references described any potential inadequacies that might have led to the incident. In contrast, the Mars Program Independent Assessment Team that was supported by NASA made approximately fifty references to the role played by headquarters [569].

The findings from the Independent Assessment Team can again be compared with the root causes that have been identified using tier analysis. Such a comparison reflects some of the limitations of this approach when applied to the less direct causes of an incident or accident. The following excerpts summarise the results of the independent enquiry:

> " Through a combination of perceived NASA Headquarters mandates and concern for loss of business, JPL and LMA committed to overly challenging programmatic goals. The JPL management perception was that no cost increase was permissible and the aggressive pricing strategy adopted by LMA exacerbated the problem... NASA Headquarters thought it was articulating program objectives, mission requirements, and constraints. JPL management was hearing these as non-negotiable program mandates (e.g., as dictated launch vehicle, specific costs and schedules, and performance requirements)... The result was that JPL management did not convey an adequate risk assessment to NASA Headquarters. What NASA Headquarters heard was JPL agreeing with and accepting objectives, requirements, and constraints. This communication dynamic prevented open and effective discussion of problems and issues. JPL management did not effectively express their concerns to NASA Headquarters about programmatic constraints, and NASA Headquarters did not seem receptive to receiving bad news... In this case, JPL and NASA

> Headquarters communications were inadequate, in part because JPL was concerned that Headquarters would perceive JPL concerns about programmatic constraints negatively; JPL did not want to antagonise the customer. NASA Headquarters was rigid in adhering to unrealistic constraints. Communication between JPL and NASA Headquarters was impeded by a cumbersome and poorly defined organisational structure within the Office of Space Science." [569]

Our use of tier analysis did not reveal many of the causal factors that are identified in the Mars Program Independent Assessment Team's report. For instance, the previous tables did not identify the communications problems that led JPL to interpret Headquarter's objectives as non-negotiable program mandates. On the other hand, the tier analysis associated a failure to encourage the use of Incident, Surprises and Anomaly reporting with Headquarters management. A number of different explanations can be proposed for such apparent differences. The first is that the subjective nature of root cause analysis, even when supported by ECF charts and tier analysis, makes it likely that different teams of investigators will focus on different aspects of an incident. It is hardly surprising, given the content of this book, that our analysis should have identified the failure of the reporting system as a root cause! A second potential explanation for these apparent differences is that the results of the tier analysis are strongly influenced by the use of ECF charts during the initial stages of an investigation. This technique encourages analysts to focus on particular events rather than on the organisational factors that create the conditions for an incident. It is important to remember, however, that this initial focus is broadened by barrier and change analysis. Both of these techniques help to ensure that causal factor analysis does look beyond the immediate events that contribute to an incident. A third explanation for the differences between the products of our tier analysis and the organisational analysis of the independent assessment team is that each of these investigations had different objectives. Our intention in identifying the root causes of the Climate Orbiter incident was to demonstrate that tier analysis could be used to identify root causes at different levels of management in each of the organisations that were involved in the incident. In contrast, the Mars Program Independent Assessment Team was more narrowly focussed on the structure and organisation of NASA's Mars Program. It therefore provides only a cursory examination of the direct events leading to the failure and certainly does not approach the level of detail shown in previous ECF charts.

The previous paragraphs have shown tier analysis can be used to identify root causes amongst the conditions and events that are derived from a causal factor analysis. An important strength of this approach is that it focuses the investigators attention on the higher levels of management within the organisations that are involved in an incident. Tier analysis also helps to explicitly distinguish generic causes, i.e., factors that might result in future failures, from the more specific causal factors that characterise a particular incident. Previous paragraphs have also identified a number of potential weaknesses. Tier analysis may be unnecessarily restrictive if it relies on causal factor analysis as a means of identifying potential causal factors. Unless this technique is used in conjunction with a broad ranging change or barrier analysis then it can be difficult to identify all of the ways in which organisational factors might contribute to an incident. Tier analysis also relies entirely upon the subjective skill of the investigator. It is possible to annotate tier diagrams in a flexible manner but they must be supported by prose descriptions if other investigators are to understand the detailed justification for identifying particular root causes from a mass of other causal factors. These descriptions are important because without them it will be difficult to validate the output from any tier analysis.

10.3.4 Non-Compliance Analysis

Rather than repeat our application of tier analysis for the Mars Polar Lander incident, this section presents an alternative form of root cause analysis. Non-compliance classification focuses on three different forms of non-compliance. The first relates to situations in which individuals *don't know* that they are violating an accepted rule or procedure. This occurs if workers receive inadequate training or if they are not informed about changes in applicable regulations. The second classification deals with situations in which individuals and teams *can't comply.* This occurs if operators or managers are denied the necessary resources to meet their obligations. The final classification relates to situations in which there is a decision not to follow rules and procedures. Individuals and teams may explicitly or implicitly decide that they *won't comply* with an applicable regulation. Table 10.22 summarises the more detailed categories that investigators must consider for each of these possible situations [207].

Don't Know:	
Never Knew	Poor training or a failure to disseminate regulations to the appropriate recipients.
Forgot	Individual factors, inadequate reminders or unrealistic assumptions on the part of an organisation about what can be recalled, especially under stress.
Didn't understand	Lack of experience or of guidance in how to apply information that has already been provided.

Can't Comply:	
Scarce Resources	Often used to excuse non-compliance. Investigators must be certain that adequate resources were requested.
Impossible	Organisations may impose contradictory constraints so that it is impossible to satisfy one regulation without breaking another.

Won't Comply:	
No penalty or no reward	There may be no incentive to comply with a requirement and hence there may be a tendency to ignore it.
Disagree	Individuals and groups may not recognise the importance of a requirement and so may refuse to satisfy it. Local knowledge may suggest that a regulation threatens safety.

Table 10.22: Root Cause Taxonomy within Non-Compliance Analysis.

The US Department of Energy recommends non-compliance analysis as a means of extracting root causes from the mass of more general causal factors that are derived from causal factor analysis [207]. The causal events that are identified using the counterfactual analysis of previous sections are associated with one of the categories shown in Table 10.22. It is worth recalling that causal factors are distinguished using the counterfactual question; would the incident have occurred if this event or condition had not held? Root causes satisfy the additional condition that they must represent a more general cause of future failures. Non-compliance analysis can be used to distinguish root causes from causal factors because each of the categories in Table 10.22 corresponds to a pre-defined set of more general root causes. By classifying a causal factor according to one of these categories, investigators are encouraged to recognise the wider problems that may stem from the associated root causes. Causal factors that fall into the *don't know* class represents a failure in the training and selection of employees. The *can't comply* class represents root causes that stem from resource allocation issues. Causal factors associated with the *won't comply* class represents a managerial failure to communicate safety objectives. For example, previous sections have used causal factor analysis to identify a number of causal factors that may have contributed to the loss of the Climate Orbiter. These included the observation that Ground-based software uses Imperial and not Metric units for thruster performance during calculation of the AMD data file. The programmers failed to follow the recommended practices that were outlined in the Software Interface Specification. Non-compliance analysis might, therefore, conclude that the software engineers never knew about this document, that they did know about it but forgot to use it or that they did not understand its relevance to the development of mission critical software. These classifications all refer to an underlying root cause; employees were not adequately trained to recognise the importance of such documents. In consequence, any remedial actions should not focus simply on the Software Interface Specification but on the more general need to ensure that software engineers have an adequate understanding of the development practices that are outlined in this and similar documents.

This approach offers a number of potential benefits for organisations whose activities are governed by well-documented guidelines, standards and regulations. Some of these documents even provide investigators with advice about how to detect the symptoms of non-compliance. For example, JPL produced a series of documents on NASA recommended practices that explicitly state what might happen if projects fail to follow the guidelines:

> "*Impact of Non-Practice:* The performance of the delivered product may be compromised if the hardware imposed limitations are not evaluated early in the design phase. Once the hardware is delivered, it is too late to select an alternative radio architecture, and there are few opportunities to mitigate the impact of any constraints on radio performance. Lacking insight into RF hardware characteristics, test engineers may waste valuable engineering hours determining the basis for the variance between expected and observed performance. For flight projects, costly problem/failure reports and project waivers will likely be processed due to the lack of an early understanding of hardware limitations." [578]

There are, however, a number of practical problems that complicate the use of non-compliance analysis as a means of identifying more general root causes from the causal factors that are identified during a causal factor analysis. Firstly, the more general root causes that are associated with the categories in Table 10.22 cannot hope to cover all of the potential root causes of adverse incidents in many different industries. in contrast, this form of analysis directs the investigators' attention towards a very limited set of factors associated with training, with resource allocation and with the communication of safety priorities. This direction can either be seen as a useful heuristic that helps to ensure consistency between analysts or as a dangerous form of bias that may obscure other underlying root causes.

The application of non-compliance analysis is further complicated by the difficulty of determining whether or not particular regulations and policy documents are applicable to particular projects. This might seem to be a trivial task in many industries. However, NASA preferred practice procedures were drafted by individual centres during the period preceding the loss of the Polar Lander and the Climate Orbiter. For example, Practice No. 1437 on end-to-end compatibility and mission simulation testing explicitly states that "all flight programs managed by the Goddard Space Flight Center (GSFC) are required to use this practice" [567]. This situation is not uncommon. Different regional or function divisions often draft supplementary regulations to support their particular activities. Problems arise when investigators must determine whether local regulations affected the course of an incident and whether they interacted with the requirements that are imposed at other levels within an organisation or from regulatory organisations.

The individual nature of many NASA projects can prevent investigators from establishing the norms that govern development and operation practices. Each project is so different that it can be difficult to identify which of those differences actually contributed to an incident. This makes it difficult for investigators to use techniques, including change analysis, that focus on the differences between 'normal' and observed behaviour. Non-compliance analysis suffers from similar problems. Differences between projects force managers to adapt existing working practices. For instance, radical changes in the relationships between JPL, NASA Headquarters and the subcontractor organisations forced program managers to adapt existing reporting procedures during the Mars Surveyor'98 program. They also complicate any attempts to enumerate those policies and regulations that govern each stage of the missions within each of the participant organisations. NASA recognise the need for flexibility in the face of changing mission demands. For instance, NASA Standard 8729.1 is one of several guidelines that specifically allows departures from the recommended practice. Such flexibility creates difficulties for investigators who must determine whether or not it was reasonable for projects to decide not to comply with recommended practice:

> "*Section 1.3 Approval of Departures from this Standard.* This standard provides guidance and is not intended for use as a mandatory requirement; therefore, there is no approval required for departing from this standard. However, the fundamental principles related to designing-in Reliability and Maintainability (R&M), as described in this standard, are considered an integral part of the systems engineering process and the ultimate R&M performance of the program/project is subject to assessment during each of the program/project subprocesses (Formulation, Approval, Implementation, and Evaluation).

A third factor that complicates non-compliance analysis is that there may be genuine uncertainty within an organisation about whether or not an individual should have complied with particular regulations. This is apparent in JPL's response to the Faster, Better, Cheaper strategy. This initiative led individual managers to reassess whether or not particular policies, for instance concerning the use of model-based validation rather than destructive testing, were still appropriate to the new context of operation:

> "(NASA, JPL and LMA) have not documented the policies and procedures that make up their FBC approach; therefore, the process is not repeatable. Rather, project managers have their own and sometimes different interpretations. This can result in missing important steps and keeping lessons learned from others who could benefit from them. [569]"

It is relatively easy in retrospect to argue that an incident occurred, therefore, a regulation was violated. It is less easy to determine whether any individuals within the organisation would have concurred with that analysis *before* the incident took place. This hindsight bias is a particular danger where non-compliance analysis is (ab)used as a mechanism for blame attribution.

It can also be difficult to apply compliance analysis to the results from previous stages in a causal factor analysis. For instance, the following list enumerate the causal factors that were identified for the Deep Space 2 and Polar Lander mishaps. These causal factors were derived by applying counterfactual reasoning to each of the events that was represented within previous ECF charts of this incident:

1. Both DS2 probes suffer electrical failure at impact
2. Forces at impact compromise aft body battery assembly
3. Forces at impact compromise RF components
4. Premature Shut-Down of engines (40 meters above surface)
5. Software marks individual legs as failed if they show spurious signals but does not reset touchdown indicator at 40 meters (entry +5:16)
6. Transient signals possible from Hall Effect magnets when legs first deploy at 1,500 meters (Entry +4:13)
7. Launch approved
8. Preliminary Design Review passed

It is difficult to directly apply non-compliance analysis to any of these causal factors. For example, the electrical failure of the Deep Space 2 probes on impact cannot itself be blamed upon a lack of knowledge about applicable regulations or on an inability to meet those regulations or on a deliberate failure to follow those regulations. This is because the causal factor related to a direct failure rather than to any particular form of non-compliance by an identifiable individual or group. A further stage of analysis is required before investigators can exploit this categorisation as a means of identifying potential root causes. For instance, the failure of Radio Frequency components on impact with the planet surface is a probable failure mode because development impact tests were limited to brassboard and breadboard components and subassemblies [579]. Visual inspections were conducted after these test to ensure that the component mountings and the electrical connections remained intact. Unfortunately, many of the components were not electrically functional. As a result, it was only possible to conduct limited inspections of the powered circuits before and after the impact tests. In other words, the impact tests established the structural integrity of the design but did not establish the functional validity. It can, therefore, be argued that the RF testing during the development of the Polar Lander indicates non-compliance with NASA requirements. In particular, Preferred Reliability Practice PT-TE-1435 governed the verification of RF hardware within JPL from February 1996. Impact tests are implied by a requirement to evaluate RF subsystem performance under 'other environmental conditions':

> "Analyses are performed early in the design of radio frequency (RF) hardware to determine hardware imposed limitations which affect radio performance. These limitations include distortion, bandwidth constraints, transfer function non-linearity, non-zero rise and fall transition time, and signal-to-noise ratio (SNR) degradation. The effects of these hardware performance impediments are measured and recorded. Performance evaluation is a reliability concern because RF hardware performance is sensitive to thermal and other environmental conditions, and reliability testing is constrained by RF temperature limitations." [578]

The failure to follow PT-TE-1435 is classified as an inability to comply. It is, therefore, associated with root causes that centre on resource allocation issues. This judgement is supported by the finding that there were several design changes late in the development program that prevented impact testing without jeopardising the launch of the Polar Lander. If the battery cells and RF subsystem assemblies had been available earlier in the development cycle then it might have been possible to comply with PT-TE-1435. This line of analysis is summarised by the non-compliance diagram illustrated in Table 10.23.

Causal Factor	Procedure or Regulation	Compliance Failure?
Forces at impact compromise RF components	Preferred Reliability Practice PT-TE-1435 Early validation of RF reliability under thermal and other environmental conditions.	Can't comply RF assembly unavailable for impact testing as design changes delay development.

Table 10.23: Non-Compliance Analysis of RF Failure Mode on Deep Space 2 Probe.

If we continue this non-compliance analysis, the situation is shown to be considerably more complex than that suggested in Table 10.23. In particular, the Preferred Practice proposed in PT-TE-1435 centres on the use of modelling as a means of validating the initial design of RF components. This is particularly important because mathematical analysis can be used to identify potential design weaknesses before projects accept the costs associated with procuring particular subsystems. PT-TE-1435 argues that these models help in situations where it is "difficult to pinpoint the exact cause of unexpected test results once the subsystem has been integrated". [578] From this it follows that the development team could have complied with PT-TE-1435 even though design changes meant that the flight unit was not available for impact tests. Mathematical models could have been used to provide the validation recommended in this regulation. Unfortunately, the impact analysis of high gravitational forces does not yield reliable results. Finite element analysis was used to validate the antenna structure. This did not provide reliable results because the impact loads were not well understood. Several antenna masts were slightly bent during impact testing, but no analytic models could be made to match the empirical damage. Empirical impact testing provides the only reliable verification method.

As before, further analysis of this apparent non-compliance can yield further insights into the complexities that characterised the development and testing of the Deep Space 2 probes. NASA requirements, such as PT-TE-1435, were well understood by JPL employees and the contractor organisations. The design changes to the RF system meant that any impact tests would not be completed before the scheduled launch of the Polar Lander. They, therefore, attempted to gain explicit approval for the decision to proceed to launch without an RF subsystem impact test:

> "The DS2 project thought there was no alternative to accepting the absence of a flight-like RF Subsystem impact test, short of missing the MPL launch opportunity. The rationale for proceeding to launch was presented and accepted at two peer reviews and presented at three project-level reviews: Risk Assessment, Mission Readiness, and Delta Mission Readiness. The project had proceed to launch concurrence from JPL and NASA upper management." [579]

Such actions can be interpreted as an understandable reluctance to comply with the requirements and recommended practices that governed RF validation. Mission schedule was interpreted within the Faster, Better, Cheaper strategy as being more critical than additional reliability tests for components that had already been validated at a structural and component level. Table 10.24, therefore, builds upon the previous analysis to document these additional reasons for non-compliance.

Causal Factor	Procedure or Regulation	Compliance Failure?
Forces at impact compromise RF components	Preferred Reliability Practice PT-TE-1435 Early validation of RF reliability under thermal and other environmental conditions.	Can't comply 1. RF assembly unavailable for impact testing as design changes delay development. 2. Mathematical modelling of high g impacts yields unreliable results. Won't comply 1. JPL and NASA upper management approve launch without RF impact validation in order for DS2 to meet launch schedule. 2. RF subsystem components had been structurally tested and were similar to other components used in previous missions.

Table 10.24: Non-Compliance Analysis of RF Failure Mode on Deep Space 2 Probe (2).

The initial resource allocation problems, connected with late design changes to RF components, were compounded by the pressures to launch on schedule. Higher-levels of management were prepared to concur with this decision, arguably, because of the perceived need to implement the the Faster, Better, Cheaper strategy. This illustrates the way in which non-compliance analysis helps to identify the deeper root causes of an incident. The specific causal factor revealed by the causal factor analysis is unlikely to threaten future missions simply because it has been identified as a potential cause of the Deep Space 2 mishap. The validation of RF assemblies will include system-level impact tests. In contrast, the root cause of the non-compliance remains a concern for subsequent missions.

Mission deadlines and tight launch schedules will continue to encourage engineers and managers to sanction non-compliance with accepted working practices. The mishap report into the management structures that contributed to the loss of the Climate Orbiter observed that:

> "NASA currently has a significant infrastructure of processes and requirements in place to enable robust program and project management, beginning with the capstone document: NASA Procedures and Guidelines 7120.5. To illustrate the sheer volume of these processes and requirements, a partial listing is provided in Appendix D. Many of these clearly have a direct bearing on mission success. This Boards review of recent project failures and successes raises questions concerning the implementation and adequacy of existing processes and requirements. If NASA programs and projects had implemented these processes in a disciplined manner, we might not have had the number of mission failures that have occurred in the recent past." [569]

The Appendix of the report lists over fifty NASA standards that were identified as relevant to this incident. These ranged from standards relating to electrical discharge control through safety-critical software development to standards for oxygen systems. This not only reflects the complexity of any non-compliance analysis, mentioned above, but it also illustrates the demands that are place on managers and operators who must ensure compliance to these regulations while also satisfying high-level mission objectives such as those implied by the Faster, Better, Cheaper strategy.

10.4 Summary

This chapter has shown how a range of diverse analytical techniques can be used to identify the causal factors that contribute to a particular incident. These causal factors can then be used to determine the underlying root causes that might continue to threaten the safety of future systems. The techniques that we have exploited are based on those advocated by the US Department of Energy. Their approach was specifically developed to support the analysis of workplace injuries. It has not been widely applied to reason about the causes of complex, technological failures. This is surprising given that NASA's Procedures and Guidelines document NPG:8621.1 on mishap reporting recommends this same approach to root cause analysis. We, therefore, demonstrated that these techniques could be used to support an investigation into the loss of the Mars Climate Orbiter and the Mars Polar Lander missions. These case studies are not 'safety-critical' in the sense that they did not threaten human life after they had left the Earth's orbit. They do, however, reflect a more general class of mission-critical incidents that are considered by many reporting systems. These case studies were also chosen because they provide an extreme example of the technological complexity and coupling that characterises many safety-critical failures. The Climate Orbiter and Polar Lander missions also provide a strong contrast with the level of technology involved in the Allentown explosion in Chapter 9.

This chapter began with the construction of ECF charts. These graphs help to identify the events and conditions that lead to an incident. They are similar to modelling techniques, especially graphical time-lines and Fault Trees, that have been introduced in previous chapters. They do, however, suffer from a number of potential limitations. In particular, ECF charts can bias investigators towards the representation of observable events rather than the wider contextual factors that made those events more likely. The US Department of Energy guidelines and the NASA procedures advocate the use of supplementary analytical techniques to uncover these factors. For instance, change analysis can be used to identify the impact that different management priorities, new working practices and technological innovation have upon the course of an incident. These changes often lead to the unanticipated interactions that have been identified as important causes of 'systemic' failures [486]. Similarly, barrier analysis helps to move the focus away from events that actively contribute to an incident. This technique encourages investigators to consider the ways in which a wide variety of potential barriers must fail in order for an incident to occur. Both of these analytical techniques can be used to look beyond the initial events that are represented in an ECF chart. They encourage investigators to revise those diagrams and, in particular, to incorporate a wider range of causal factors.

The causal factors are distinguished from a wider range of contextual factors using causal factor analysis. This technique involves the use of counterfactual reasoning. For each event in the revised ECF chart, investigators must ask 'would the incident have occurred without this event?'. If the answer is yes then the event is not considered to be a causal factor. If the answer is no then investigators must record further information about the event. This information centres on a number of prompts including: what led to the event? What went wrong? How did the barriers fail? Who was involved in the event? Is the event linked to a more general deficiency? The results of this more detailed analysis can be recorded in an ECF table. These, in turn, are used to drive any subsequent root cause analysis.

Causal factors are identified using counterfactual reasoning. An incident would not have occurred, if the event or condition had not occurred. In contrast, root causes are events or conditions that threaten the safety of future systems. They often result from the amalgamation of several causal factors. For example, the failure of several barriers may indicate a more general failure to ensure adequate protection. Any attempt to fix particular barriers will still leave a concern that other barriers may still be susceptible to other forms of failure until this root cause is more directly addressed. Several techniques have been proposed to help investigators move from specific causal factors to these more general root causes. Again our use of tier and non-compliance analysis has been guided by the US Department of Energy's recommendation. Tier analysis depends upon the development of tables that associate causal factors with different levels in an organisational structure. The entries in these tables are then inspected in order to identify more general patterns that might indicate a root cause that is common to several causal factors. In contrast, non-compliance analysis involves the examination of any rules or procedures that might have been violated either directly by an event or by the wider conditions that made an event more likely.

It is important to emphasis that the techniques which we have described do not provide a panacea for the problems of root cause analysis. It can be difficult to apply some of these approaches to the specific circumstances that characterise particular technological failures. The documentation techniques that are associated with key stages in the analysis, especially the revised ECF charts, are cumbersome and intractable. All of the techniques that we have described rely upon the subjective skill and experience of individual investigators. The insights that they provide must, therefore, be validated by other members of an investigation team or a safety managementgroup. A number of researchers are currently working to produce automated systems that remove some of the subjectivity involved in root cause analysis. Unfortunately, sophisticated reasoning tools often impose unacceptable constraints upon the way in which an incident is modelled. The syntax and semantics of any input must be narrowly defined so that the system can recognise and manipulate model components during any subsequent root cause analysis. There are a number of potential solutions to this problem, including structural induction over graphical structures similar to ECF chart. In anticipation of the results of this research, it is difficult to underestimate the importance of the tables and diagrams that are presented in this chapter. They provide other analysts and investigators with means of tracing the reasons why particular events and conditions are identified as causal factors. They also help to document the process by which root causes are determined. Without such documents, it would be extremely difficult to validate the subjective analysis of incident investigators.

The penultimate remarks in the Chapter belong to Daniel Goldin; the NASA Administrator who first formulated the Faster, Better, Cheaper strategy. He spoke to the engineers and managers at the Jet Propulsion Laboratory about the loss of the Climate Orbiter and the Polar Lander.

> "I told them that in my effort to empower people, I pushed too hard... and in so doing, stretched the system too thin. It wasn't intentional. It wasn't malicious. I believed in the vision... but it may have made failure inevitable. I wanted to demonstrate to the world that we could do things much better than anyone else. And you delivered – you delivered with Mars Pathfinder... With Mars Global Surveyor... With Deep Space 1. We pushed the boundaries like never before... and had not yet reached what we thought was the limit. Not until Mars 98. I salute that team's courage and conviction. And make no mistake: they need not apologise to anyone. They did not fail alone. As the head of NASA, I accept the responsibility. If anything, the system failed them." [574]

There is a danger that the recent emphasis on systemic failures will discourage investigators from pursuing the coherent analysis of specific root causes. Many incidents are characterised by emergent behaviours that stem from complex interactions between management practices, operational procedures and particular technologies. These interactions are not, however, random. They are shaped and directed by the regulatory environment and by higher-levels of management. Goldin's words are important because they acknowledge personal and corporate responsibility for the systemic factors that led to failure.

Chapter 11

Alternative Causal Analysis Techniques

The previous chapter showed how a range of existing techniques can be applied to identify the root causes and causal factors that lead to failures in high-technology systems. In particular, we have shown how Events and Causal Factor (ECF) charts can be derived from the findings of primary and secondary investigations. The scope of these diagrams can be both broadened and deepened using barrier analysis and change analysis. Counterfactual reasoning can then be applied to distinguish causal factors from other contextual influences on an incident. Finally, tier analysis and non-compliance analysis can be used to distinguish the root causes that threaten future safety from the causal factors that characterise individual incidents. The intention was to provide a relatively detailed case study in the application of these particular analytical techniques. The choice of approach was motivated by the recommendations of the US Department of Energy and of NASA NPG 8621.1.

The following pages build on this analysis by introducing a range of alternative techniques. The intention is to provide a broader perspective on causal analysis. As we shall see, some of these techniques can be integrated into the approach that was described in the previous chapter. For instance, ECF charts can be replaced by Sequentially Timed and Events Plotting or by Multilinear Events Sequencing [72, 346]. The justification for broadening the scope of the previous chapters is that there have been few investigations into the comparative utility of causal analysis techniques. There are some notable exceptions. For instance, Benner [73] provides a rating of accident models and investigative methods. Munson has more recently presented a comparative analysis of accident modelling techniques applied to Wildland Fire Fighting Incidents [552]. It is important to note that both of these studies are more concerned with the range of factors that are captured by particular modelling notations and their integration into investigatory processes. Neither directly studies the ultimate application of these models to support causal analysis. In the absence of such comparative studies, it is important that investigators have a clear idea of the alternative approaches that might be used to support the causal analysis of safety-critical incidents.

A fire on-board the bulk carrier Nanticoke provides a case study for the remainder of this chapter [621] This is appropriate because it provides a further contrast to the Pipeline expolosion that was modelled in Chapter 9 and the Mars case studies that were analysed in Chapter 10. The Nanticoke departed Camden, New Jersey, USA, on 19 July 1999. It was carrying 29,000 tons of petroleum coke. Between 12:00-16:00 on the 20th July, an engineer cleaned the forward fuel filter on the Nanticoke's port generator as part of a preventive maintenance routine. The engineer started the generator and tested the filter for leaks around 15:00. At 15:15 the chief engineer entered the engine-room and inspected the generators. He found that all temperatures and pressures were normal and, therefore, continued on to the control room. Shortly after this, a fire drill was started. The chief engineer relieved the duty engineer who had to go to an assigned fire station. During this time, the chief engineer and a mechanical assistant maintained their watch from the engine control room where they could not directly observe the state of the generator. The fire drill ended at 16:00. Shortly after this, the chief engineer noted a high cooling water temperature alarm from port generator cylinder No. 1 from the engine control room displays. He left the control room and discovered that the engine-room was full of smoke.

The chief engineer returned to the control room and sounded the general alarm. He then called the bridge and informed them of the fire. He shut down the port generator, isolated its fuel supply and then put on a smoke hood so that he could find the mechanical assistant. The mechanical assistant had already left the engine-room and so the chief engineer returned to the control room. The control room was not equipped with an emergency exit and so he was forced to follow handrails to the engine-room exit door on the main deck. The starboard generator was left running to supply power to the rest of the vessel.

On the bridge, the master sent a security call that was acknowledged by the United States Coast Guard in New York City. They then transmitted a Mayday as the extent of the fire became more apparent. The fire parties were standing down from the drill and were in the process of removing their protective fire suits when the alarm sounded. Two crew members entered the engine-room using air packs and protective suits that were already to-hand following the fire drill. They initially used carbon dioxide extinguishers to fight the fire but were driven back by the heat. A second team then repeated the attempt using a fire hose but this also failed to completely extinguish the fire. The chief engineer then performed a headcount and ensured that the engine-room vents were closed. He then discharged the Halon extinguishing system around 16:40. The fire was fully extinguished by 17:22. Shortly after this time, the gangway doors were opened to ventilate the engine-room.

The remaining pages use this incident as a case study to illustrate a number of alternative causal analysis techniques. This provides investigators with an overview of the rival approaches to the ECF and Causal Analysis techniques that were presented in Chapter 10. The following pages also introduce complementary techniques that can be used to supplement or replace the method that was described in the previous chapter.

11.1 Event-Based Approaches

ECF charts can be used to analyse the way in which various chains of events and conditions contribute to safety-critical incidents. Failure sequences can be sketched, edited and extended as other techniques, such as barrier analysis, drive further insights into an incident. Unfortunately, a number of limitations reduce the utility of this approach. For instance, Munson argues this method is labour intensive and often requires considerable amounts of time to complete even a preliminary analysis [552]. It also requires a considerable range of domain knowledge, in additional to the technical knowledge required to perform the analysis [290]. For instance, tier analysis relies upon a knowledge of the managerial structure of the many organisations that are involved in an incident. As we have seen in the previous chapter, commercial barriers and the complexity of some management organisations can frustrate attempts to elicit this information even in cases where serious failures have occurred. Further limitations stem from the manner in which temporal information is included within individual events and conditions. There is an implicit assumption that time flows from the left to the right in an ECF chart. An event or condition is assumed to occur after events or conditions that are placed to their left. There is, however, no time scale associated with ECF charts. In consequence, investigators must manually search through dozens of nodes in these diagrams to determine what might have happened at any particular moment in time.

11.1.1 Multilinear Events Sequencing (MES)

Multilinear Events Sequencing (MES) provides an alternative to the ECF charts in Chapter 10. It is different from the more general modeling techniques introduced in Chapter 9, such as Petri Nets and Fault Trees, because it was specifically developed to support accident and incident analysis [72, 346]. It is intended to help investigators model and analyse an incident as an investigation progresses [705]. This implies that the approach avoids some of the overheads associated with the more elaborate techniques that are presented in previous chapters.

The basic premise that underlies MES analysis is that both successful operations and failures are the result of processes that are comprised of interactions between events. Rimson and Benner go on to argue that incidents occur "when changes during a planned process initiate an unplanned process which ends in an undesired outcome" [705]. Such comments must be balanced against situations in which two planned processes interact to produce an undesired outcome [449]. The underlying assumption, however, is that by analysing changes in a planned process it is possible to identify the potential causal factors that lead to adverse events. Processes are described in terms of a relationship between events. This is very similar to the approach adopted in ECF charts.

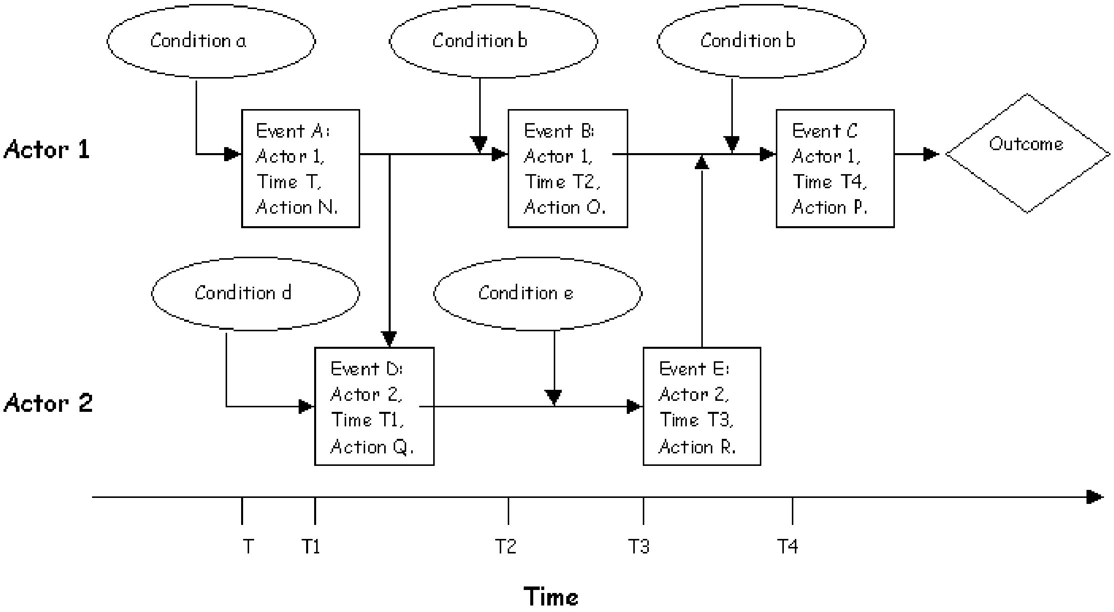
Condition a
Condition b
Condition b
Actor 1
Event A: Actor 1, Time T, Action N.
Event B: Actor 1, Time T2, Action O.
Event C Actor 1, Time T4, Action P.
Outcome
Condition d
Condition e
Actor 2
Event D: Actor 2, Time T1, Action Q.
Event E: Actor 2, Time T3, Action R.
T
T1
T2
T3
T4
Time

Figure 11.1: Abstract View of A Multilinear Events Sequence (MES) Diagram

Figure 11.1 presents the high-level form of MES flowchart. Each of the events in Figure 11.1 is described in terms of a block of information. These represent an actor performing an action at a particular time. A time-line is also included at the bottom of MES charts. This is used to explicitly represent the timing of events. It is important to note, however, that the relative position of a condition does not explicitly convey any temporal information. As can be seen, there is a deliberate attempt to help investigators identify situations in which simultaneous events contribute to an incident. The intention is to to "discover possible unknown linking events, causes, and contributing factors" [552]. The resulting diagrams resemble annotated flowcharts. This should not be surprising. The developers of MES argue that "if you can't depict a process in a flowchart, you don't understand it!" [705]. Such statements should be interpreted with care. The underlying importance of constructing accident models that are easily understood by a number of different investigators cannot, however, be denied. The MES methodology can be summarised as follows:

1. *Identify the boundaries of an incident.* A key objective behind the development of MES was to construct a method that could be used to delineate the beginning and the end of an accident sequence. Peturbation Theory (or P-Theory) was proposed to support these objectives. This starts from the assumption that the "dynamic equilibrium of successive events progresses in a state of homeostasis requiring adaptive behaviour or adaptive learning by the actors involved in maintaining the stable flow of events" [72]. Incident sequences begin with a perturbation that disturbs this dynamic equilibrium. If the system adapt to these changes then homeostasis can be maintained. If the system fails to adapt then an accident or incident sequence begins. Initial

peturbations can initiate cascading sequences of events that, in turn, place further pressures on other system components. These components can either fail or they can adapt to changing circumstances. P-Theory defines a 'near miss' incident to occur if system components adapt to any perturbations before an injury or other form of loss occurs. A number of caveats can be applied to this aspect of the MES technique. Some authors have proposed that the search for peturbations should end when "the final damaging event" is identified [552]. As we have seen, however, any analysis should ideally also consider the immediate response to any adverse occurrence given that this can either exacerbate or mitigate the consequences of any initial failure. Secondly, there are some incidents in which it is difficult ever to identify homeostasis. For instance, the relationship between LMA and JPL continued to evolve throughout the Mars Surveyor'98 missions. It is, therefore, very difficult to apply P-Theory as a means of identifying any single external event that triggered the failures. It is important to reiterate, however, that the intention behind P-Theory is simply to establish the boundaries of an incident so that investigators can begin to delineate the more detailed flow of events that contribute to a failure.

2. *Construct event blocks.* Investigators must construct a 'block' of information about each event that leads to an incident. This information must identify the actor that is associated with each event. It must also identify the action that led to the event. Both the actor and their action must be described as precisely as possible without "qualitative adjectives, adverbs, or phrases" [72]. Finally, investigators should note the time at which the event occurred. These requirements can raise a number of practical difficulties. Previous chapters have described the reliability problems that often frustrate attempts to use advanced automated logging and tracking systems to derive precise timings for critical events in the aftermath of an incident or accident. It can also be difficult to identify the agent that is associated with the ignition of the fire onboard the Nanticoke. The most probable high-temperature sources were identified as the indicator tap that protruded from the generators cylinder head and an uncovered exhaust manifold associated with the engine's turbocharger. Neither of these inanimate objects can easily be interpreted as agents even though the ignition event is critical to an understanding of the incident. One solution is to extend the notion of 'agency' to include systems and subsystems that exhibit particular behaviours in response to environmental changes. The developers of the MES method have an even broader interpretation in which actors include inanimate objects such as tires, machines and even water [72].

3. *Construct an MES flowchart.* An MES flowchart maps each event block onto two axes. The X axis represents the actors involved in an incident. In the Nanticoke case study, the master could be listed above the engineer. The engineer, in turn, might be inserted above the mechanical assistant and so on. The Y axis denotes the passage of time during an incident. The developers of the MES approach argue that because each actor is typically involved in a number of sequential events, their actions will appear as a horizontal line of event blocks across the chart. Again, this raises a number of concerns. Firstly, human factors research has shown that operators often interleave sub-tasks [666]. Interruptions can force individuals to suspend particular actions only to resume them once the immediate situation has been addressed. Similarly, it is a routine occurrence for operators to simultaneously perform multiple control tasks. Further problems stem from the construction of the MES flowchart. The granularity of the time-line must be appropriate to the circumstances that are being considered. As we have seen for time-lines, this can cause problems for incidents that are characterised by distal events that occur many months before a large number of more proximal failures. In consequence, investigators can be forced to exploit differing time-scales over the period under consideration. Each event block is then inserted into the two-dimensional array at the position determined both by the agent responsible for the event and the time at which the event is assumed to occur.

4. *Identify Conditions* The construction of an MES flowchart provides investigators with an overview of the events leading to an incident. This, in turn, can help to identify those condi-

tions that make particular events more likely to occur. Each condition is linked to at least one event using an arrow. Each condition can itself be the outcome of other external peturbations. These events can also be introduced into an MES flowchart, providing analysts with a further means of expanding the scope of any investigation. This process helps analyst to explore the underlying conditions that might trigger future perturbations and, hence, could trigger any recurrence of an incident. Experience in applying the MES approach persuaded its developers that conditions ought to be omitted from subsequent versions of the technique. It was argued that the inclusion of conditions in the MES flowcharts is superfluous because conditions are stable until changed by some action. Investigators should, therefore, focus on analysing the events that characterise an incident. This is an important difference between the version of MES that is used in this section, where conditions are included, and the STEP methodology in the following section, where conditions are omitted.

5. *Validate the assignment of event blocks within the flowchart.* After having constructed an initial flowchart, analysts must ensure that they have a coherent model of the events leading to an incident. This involves two checks. Firstly, they must ensure that the array accurately reflects the ordering for each pair of events performed by any agent. In other words, investigators must ensure that all events to the right of any particular event occur after that event. Secondly, analysts must ask whether the preceding events are both necessary and sufficiency for any following events to occur. Additional analysis must be performed if either of these tests fails. For example, the labels that are used to identify each event can be ambiguous. In such circumstances, investigators may be forced to break them down into more detailed 'sub-events'. Alternatively, events may have been omitted during the early stages of any investigation. Additional evidence can be gathered to identify any missing event blocks.

6. *Identify causal relationships.* The second of the two validation criteria, mentioned above, can be used to identify causal relationships between event blocks. Investigators annotate the flowchart so that it is possible to identify the necessary and sufficient conditions for each event to occur. Arrows can be used to represent a causal relationship between events and conditions. It should be emphasised that this is orthogonal to the temporal relationships that are denoted along the X-axis of the MES flowchart. Once this has been done, it is important that investigators consider whether there are any alternative causal hypotheses that are not reflected by the relationships that have been denoted on the flowchart. For instance, an oil leak from the forward filter cover and the ignition source provided by the turbocharger exhaust together describe sufficient conditions for the Nanticoke fire. Each event is also necessary in order for the incident to occur. There may, however, be other causal explanations. For instance, the ignition source might have been provided by the indicator tap. Either of these hypotheses might provide the necessary conditions for the subsequent mishap. Analysts must, therefore, conduct further investigations including reconstructions and empirical studies to determine which of the hypotheses is most likely. The previous requirements of temporal coherence and causal 'sufficiency' should again be applied if the chart is revised to reflect a new hypothesis. This stage is important because it encourages analysts to consider whether there may be alternative causal complexes that might have resulted in the same consequences [508]. There are further benefits. For instance, it is possible to compare the causal model in the MES flowchart with alternative models of the intended process behaviour. This can be used not only to identify the external peturbations that lead to an incident but also the ways in which internal barriers must fail in order for an incident to progress.

7. *Identify corrective actions.* Investigators can annotate the resulting MES flowchart to denote any events or conditions that should form the focus for future interventions. These potential intervention points must be analysed to identity means of mitigating the undesired outcome or of making any peturbations less likely. Recommendations can then be made to commissioning and regulatory authorities.

Figure 11.2 illustrates the results of an initial MES analysis of the Nanticoke case study. The analysis begins by identifying an event that disturbs the previous homeostasis or equilibrium of the system.

P-Theory suggests that if the system adapts to these changes then homeostasis can be maintained. If the system fails to adapt then an accident or incident sequence begins. Initial peturbations can lead to cascading sequences of events that, in turn, place further pressures on other system components. In the Nanticoke example, modifications to the forward filter removed the seating grooves that helped to maintain a seal between the copper gasket and its securing bolt. This created problems for the watchkeeping engineer when they attempted to achieve such a seal.

Figure 11.2 illustrates further events that contributed to the engineer's problems. Copper gaskets are often deformed by the pressures that they sustain under normal operating conditions in engine filters. The engineer was, however, forced to anneal and re-use the existing component as there were no spares on-board the Nanticoke. Under normal circumstances, this need not have had serious consequences. However, the deformation of the gasket may have contributed to the engineer's difficulties in sealing the filter assembly. As can be seen, the time-line in the MES flowchart provides a reference point for th events that contributed to this incident. The fuel started to escape under pressure at some point after the Chief Engineer's inspection at 15:15. The fuel was ignited by a source on the port generator at some time after it started to spray from the filter.

Figure 11.2 illustrates some of the issues that complicate the development of MES flowcharts. For instance, the ignition event, labelled D2, is associated with the engine as a whole. This diagram could, however, be refined to represent a lower level of detail. The ignition source was either the exposed indicator tap or the exhaust manifold. These two agents could be introduced to replace the generator. Unfortunately, this creates further problems. The proponents of the MES approach argue that investigators must minimise any uncertainty over the events that contribute to incidents and accidents [72]. MES flowcharts do not have any equivalent of an OR gate in a fault tree. In consequence, it is difficult to denote that the ignition source was either the indicator or the manifold. Figure 11.2 therefore refers to the port generator rather than its specific components. Part/sub-part ambiguity is used to avoid the disjunction associated with alternative events. Ideally, such imprecision might be avoided by empirical tests and simulations. As we have seen, however, it is not always easy to obtain the resources that are required to support such investigations even in the aftermath of safety-critical incidents.

Benner's P-Theory suggests that incidents are distinguished from accidents by the manner in which the system regains equilibrium without adverse consequences. This is illustrated by the outcome event in Figure 11.2. This is linked to three other events. D2 described the ignition of the fuel source. C3 describes the escalation of the fire after the O-rings on the filter's main covers were melted. Event E3, in contrast, describes the Chief Engineer's mitigating actions in shutting down the port generator.

Figure 11.3 introduces a number of conditions into the event structure that was shown in Figure 11.2. This follows the general approach that was introduced for ECF charts. The use of events and conditions offers a number of benefits. In particular, it helps to distinguish between an event and its outcome. This is illustrated by the event A1. In Figure 11.2 this was initially used to denote modification to forward filter cover/bolt sealing surface removes groove for copper washer. This captures the event, the maintenance, as well as its outcome, the removal of the seating groove. In Figure 11.3 the event is simplified to Modifies forward filter cover/bolt sealing surface. The outcome is denoted by a condition Grooves for copper washer are removed, sealing surface is uneven and grooved with file marks. These distinctions are important for the subsequent analysis of an incident. By separating the representation of an event from its outcome, analysts are encouraged to think of alternative consequences for key events during any mishap. In this instance, it may not be possible to prevent future modifications to the sealing surface but action could be taken to ensure that the sealing surface is levelled prior to operation.

Further conditions help to explain the reasons why particular events occurred. For instance, we had to explain why the watchkeeping engineer annealed the existing copper gasket, denoted by event B1 in Figure 11.2. In contrast, Figure 11.3 introduces a condition to explain that there were no spare gaskets on board the vessel at the time of the maintenance operations. A condition is also used to explain that the deformation of a gasket, event C1, can make it difficult to obtain a good seal. Finally, Figure 11.3 introduces a condition to explain that the ignition, denoted by event D2, was possible at temperatures below the flash point for the fuel because it was being sprayed under

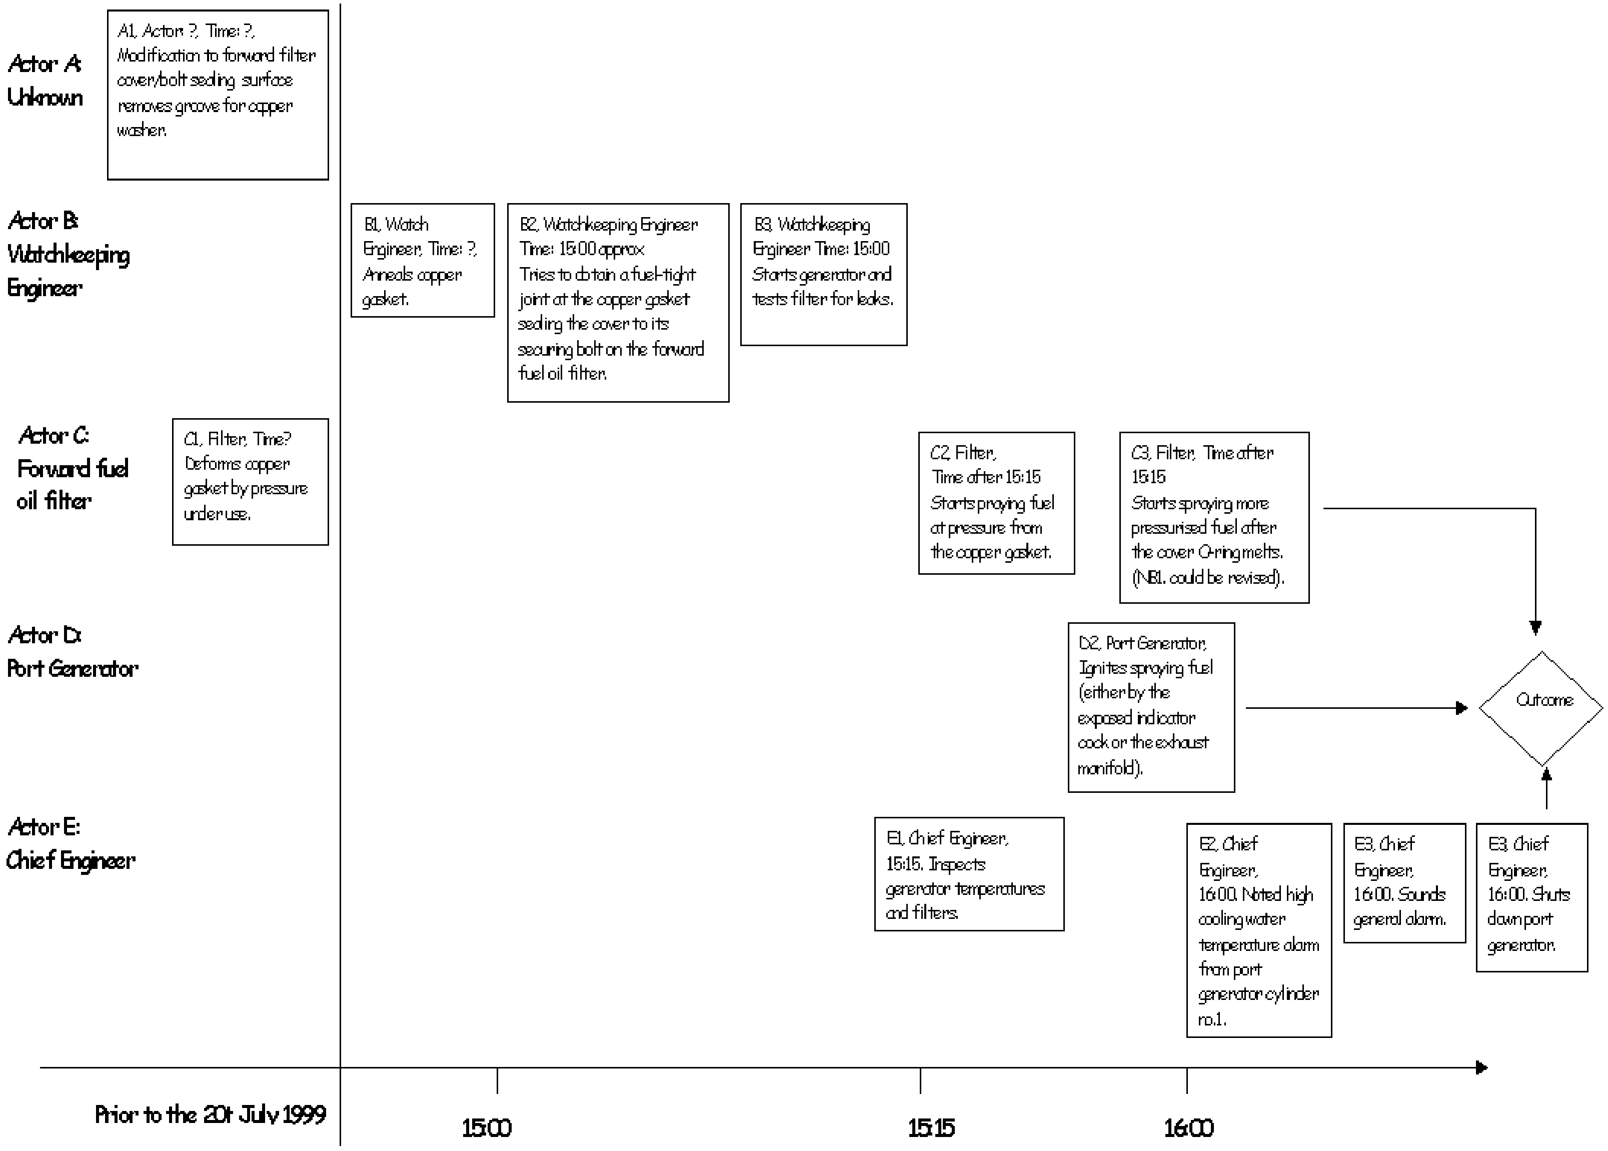
Actor A:
Unknown
A1, Actor ?, Time: ?,
Modification to forward filter cover/bolt sealing surface removes groove for copper washer.
Actor B:
Watchkeeping Engineer
B1, Watch Engineer, Time: ?,
Anneals copper gasket.
B2, Watchkeeping Engineer
Time: 15:00 approx
Tries to obtain a fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter.
B3, Watchkeeping Engineer Time: 15:00
Starts generator and tests filter for leaks.
Actor C:
Forward fuel oil filter
C1, Filter, Time?
Deforms copper gasket by pressure under use.
C2, Filter,
Time after 15:15
Starts praying fuel at pressure from the copper gasket.
C3, Filter, Time after 15:15
Starts spraying more pressurised fuel after the cover O-ring melts.
(NB1. could be revised).
Actor D:
Port Generator
D2, Port Generator,
Ignites spraying fuel (either by the exposed indicator cock or the exhaust manifold).
Outcome
Actor E:
Chief Engineer
E1, Chief Engineer,
15:15, Inspects generator temperatures and filters.
E2, Chief Engineer,
16:00, Noted high cooling water temperature alarm from port generator cylinder no.1.
E3, Chief Engineer,
16:00, Sounds general alarm.
E3, Chief Engineer,
16:00, Shuts down port generator.
Prior to the 20t July 1999
15:00
15:15
16:00

Figure 11.2: An Initial Multilinear Events Sequence (MES) Diagram

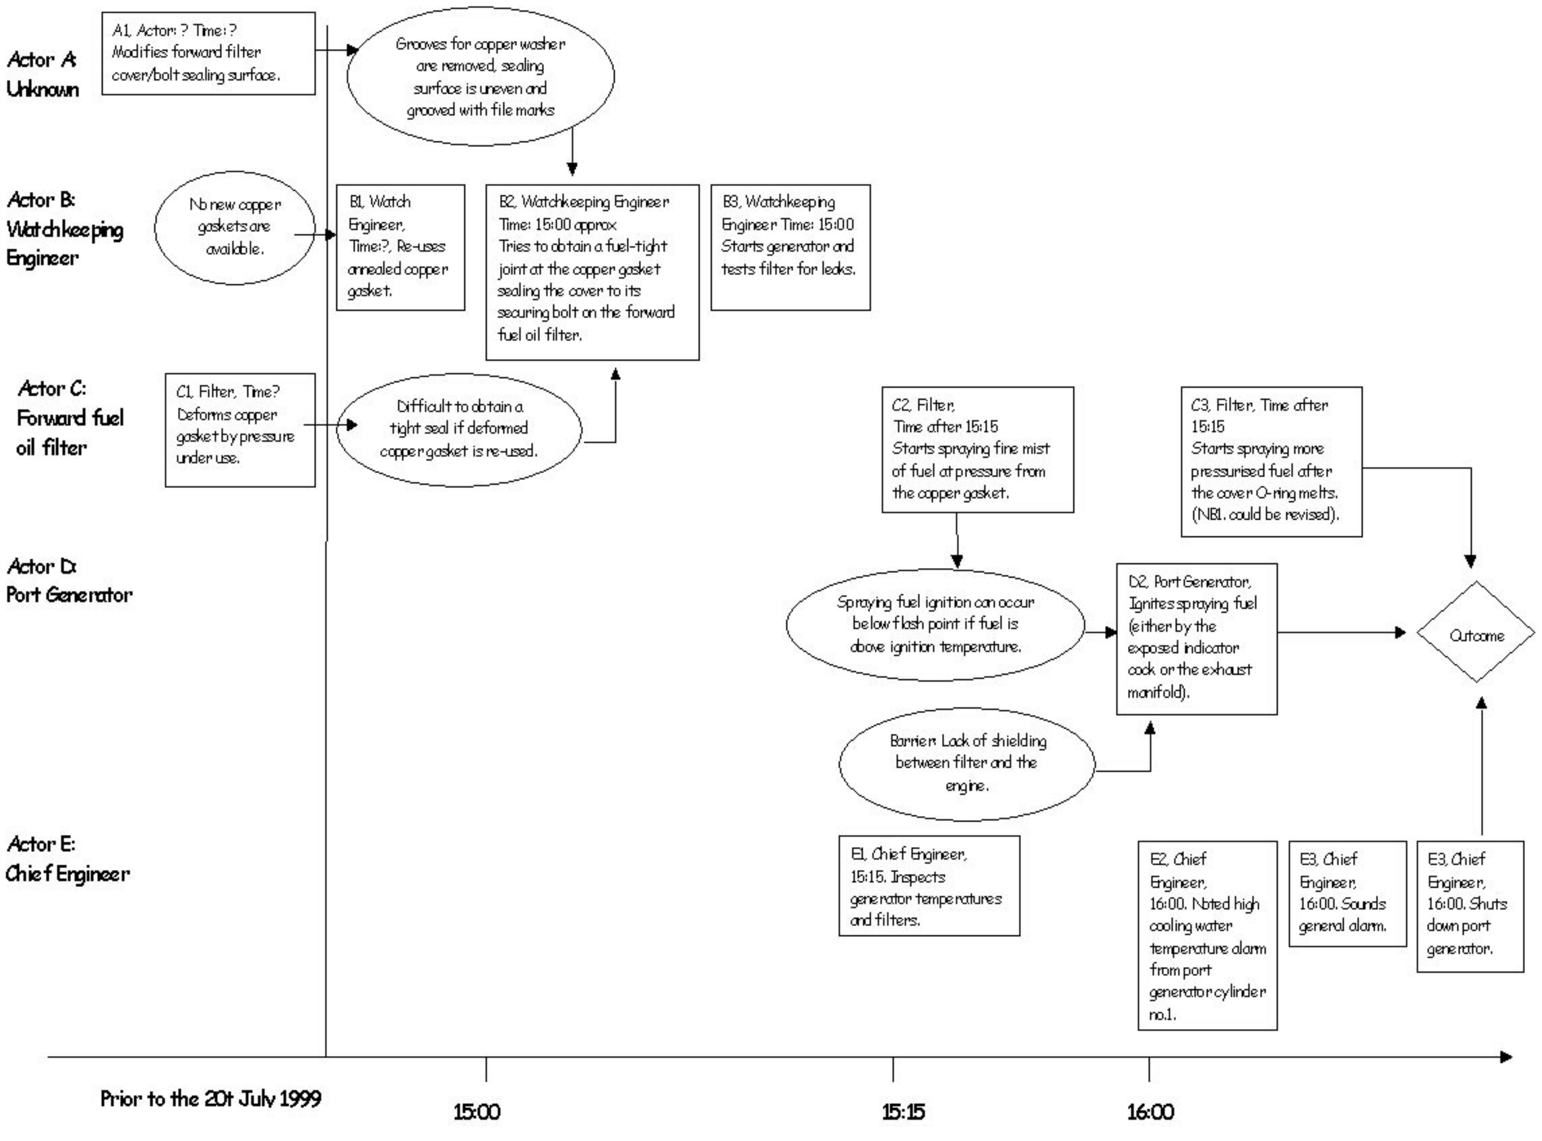
Actor A: Unknown
A1, Actor: ? Time: ? Modifies forward filter cover/bolt sealing surface.
Grooves for copper washer are removed, sealing surface is uneven and grooved with file marks
Actor B: Watchkeeping Engineer
No new copper gaskets are available.
B1, Watch Engineer, Time:?, Re-uses annealed copper gasket.
B2, Watchkeeping Engineer Time: 15:00 approx Tries to obtain a fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter.
B3, Watchkeeping Engineer Time: 15:00 Starts generator and tests filter for leaks.
Actor C: Forward fuel oil filter
C1, Filter, Time? Deforms copper gasket by pressure under use.
Difficult to obtain a tight seal if deformed copper gasket is re-used.
C2, Filter, Time after 15:15 Starts spraying fine mist of fuel at pressure from the copper gasket.
C3, Filter, Time after 15:15 Starts spraying more pressurised fuel after the cover O-ring melts. (NB1. could be revised).
Actor D: Port Generator
Spraying fuel ignition can occur below flash point if fuel is above ignition temperature.
D2, Port Generator, Ignites spraying fuel (either by the exposed indicator cock or the exhaust manifold).
Barrier: Lack of shielding between filter and the engine.
Outcome
Actor E: Chief Engineer
E1, Chief Engineer, 15:15. Inspects generator temperatures and filters.
E2, Chief Engineer, 16:00. Noted high cooling water temperature alarm from port generator cylinder no.1.
E3, Chief Engineer, 16:00. Sounds general alarm.
E3, Chief Engineer, 16:00. Shuts down port generator.
Prior to the 20t July 1999
15:00
15:15
16:00

Figure 11.3: A MES Flowchart showing Conditions in the Nanticoke Case Study

pressure, denoted by event C2. These conditions do not simply to separate out information about an event and its consequences. They provide important contextual details that can help the members of a multidisciplinary investigation team to understand the significance of particular events. The importance of this should not be underestimated. Without such explicit annotations, investigators may rely upon inappropriate assumptions about their colleagues' ability to reconstruct the ways in which particular events contribute to the course of an incident.

Figure 11.3 extends the notation described in Benner's original work [72]. A condition represents the absence of a barrier; lack of shielding between the filter and the engine. The initial MES notation makes it difficult for analysts to represent both the absence of barriers and errors of omission. This is entirely deliberate. It can be argued that investigators must focus on what *did* happen during an incident rather than what *might* have happened. By drawing other investigators' attention to the absence of particular protection measures, analysts can potentially obscure information about the performance of those barriers that were available. These objections also argue against our previous use of barrier analysis to drive ECF modelling in Chapter 9. A number of arguments support our use of conditions to represent the lack of shielding in Figure 11.3. Firstly, there is no empirical evidence to support either position in this argument. Until such evidence is obtained it is difficult to determine whether or not the introduction of information about missing barriers will bias an investigator's analysis of an incident. Secondly, even if information about errors of omission and absent barriers are excluded from incident models, these details must be explicitly considered during any subsequent analysis.

Figure 11.4 illustrates the results of introducing causal information into Figure 11.3. As mentioned above, this involves a variant of the counterfactual reasoning introduced in previous chapters. Starting with the earliest event or condition on the time-line, analysts must ask whether the next event or condition in time would have happened if this earliest event had not occurred. If the answer is no then they form a causal pair and an arrow can be drawn from the leftmost event or condition to the related event or condition. If the answer is yes then the earliest event or condition is not a necessary cause of the subsequent event or condition. No link is drawn. The analysis continues until the investigators has asked whether all of the subsequent events or conditions were potential causes by the initial event or condition. The entire process is then repeated for each subsequent event or condition in the MES flowchart. In practice, however, a number of 'optimisations' are often made. For instance, transitive links are omitted. If event or condition A causes event of condition B, which in turn, causes event or condition C then arrows need only be drawn between A and B and between B and C. The causal arrow between A and C is implied.

Causal analysis can help investigators to identify potential revisions to an existing MES flowchart. For example, Figure 11.4 introduces an event labelled fuel tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter fails. This proved necessary in order to link the previous observations about the watchkeeping engineers difficulty in obtaining a seal to the later events that described the course of the fire. As can be seen, a question mark is used to denote a degree of uncertainty in this causal link. Without it, however, there would have been no explicit means of representing that the maintenance task was a potential cause of the incident.

Figure 11.4 also illustrates the way in which a causal analysis can help to identify events that are otherwise isolated from the causal 'flow' that leads to an incident. In this case, there is an event which denotes that the Chief Engineer inspects generator temperatures and filters at 15:15. This event is important for our understanding of the incident because it helps us to determine that the fire did not take hold before that moment in time. It does not, however, play a direct role in the incident. The proponents of MES analysis, therefore, argue that it ought to be omitted from future diagrams. It is important not to underestimate the pragmatic benefits of such guidelines. It is very rare to find that any modelling or causal analysis technique provides advice about when *not* to introduce additional information that might obscure or otherwise hinder subsequent investigations.

A number of limitations can be identified with the MES techniques described in this section. As mentioned, the developers found that investigators used conditions in an arbitrary and ad hoc manner. Previous sections have argued that this is an important strength of the ECF approach. Investigators can use conditions to denote broad insights into the context in which an incident occurs. In contrast, Benner views this as a dangerous abuse because conditions can introduce superfluous

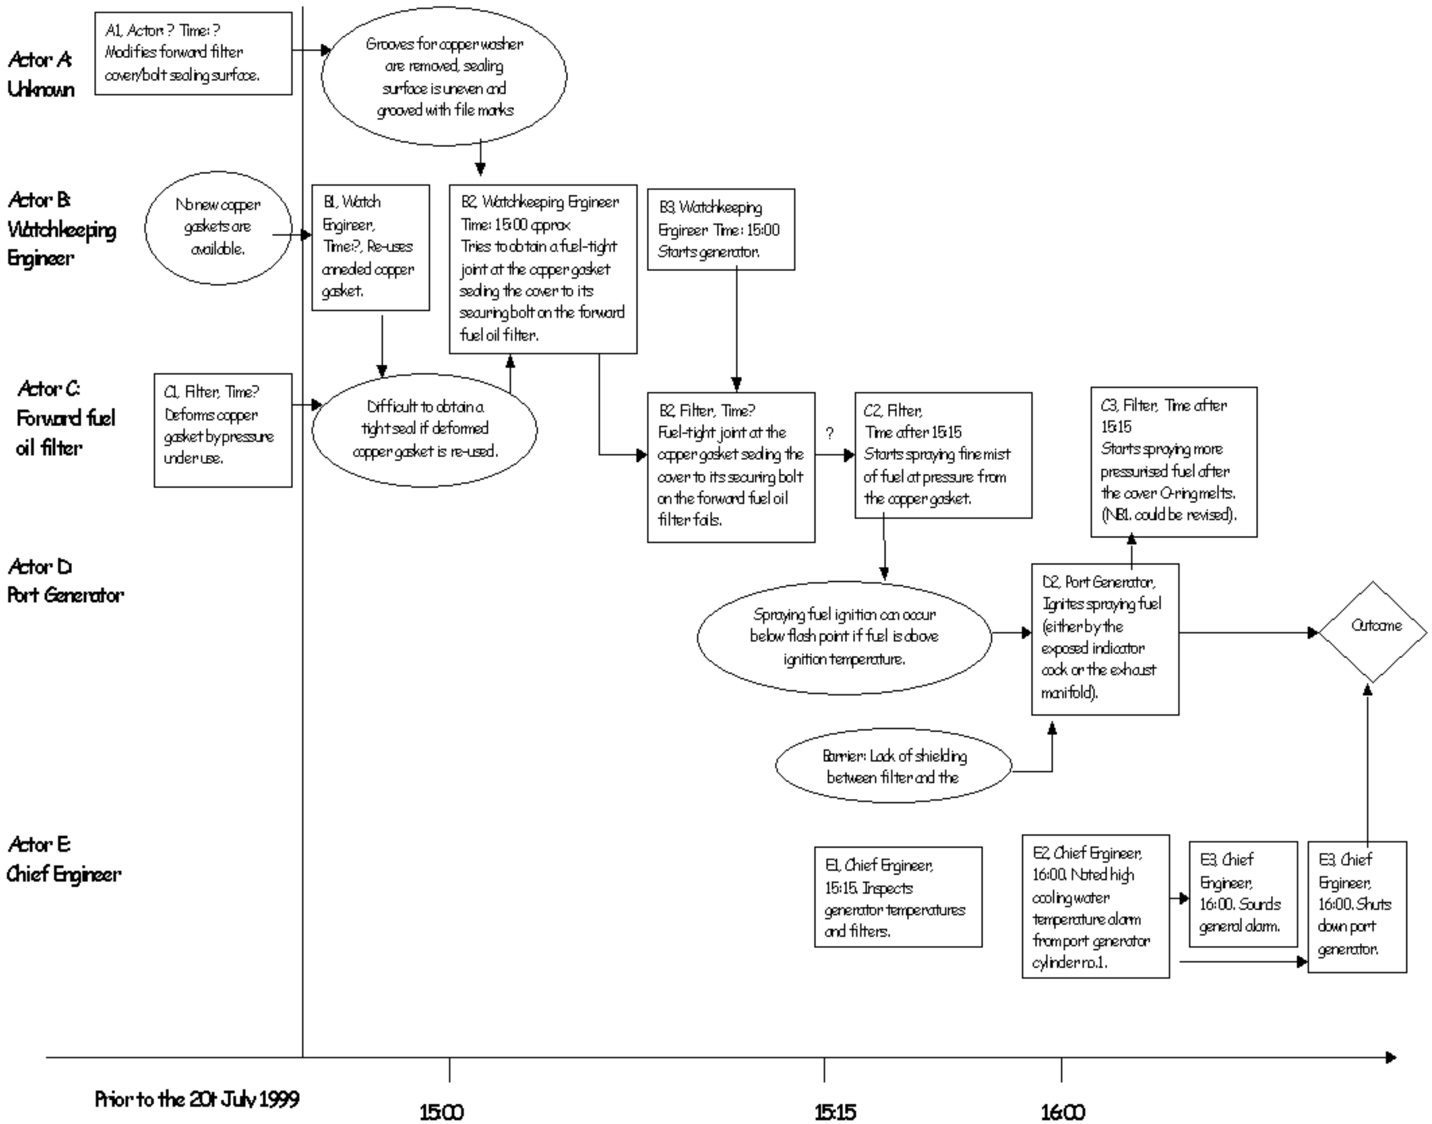
Actor A: Unknown
A1, Actor ? Time: ? Modifies forward filter cover/bolt sealing surface.
Grooves for copper washer are removed, sealing surface is uneven and grooved with file marks
Actor B: Watchkeeping Engineer
No new copper gaskets are available.
B1, Watch Engineer, Time:?, Re-uses annealed copper gasket.
B2, Watchkeeping Engineer Time: 15:00 approx Tries to obtain a fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter.
B3, Watchkeeping Engineer Time: 15:00 Starts generator.
Actor C: Forward fuel oil filter
C1, Filter, Time? Deforms copper gasket by pressure under use.
Difficult to obtain a tight seal if deformed copper gasket is re-used.
B2, Filter, Time? Fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter fails.
?
C2, Filter, Time after 15:15 Starts spraying fine mist of fuel at pressure from the copper gasket.
C3, Filter, Time after 15:15 Starts spraying more pressurised fuel after the cover O-ring melts. (NB1. could be revised).
Actor D: Port Generator
Spraying fuel ignition can occur below flash point if fuel is above ignition temperature.
D2, Port Generator, Ignites spraying fuel (either by the exposed indicator cock or the exhaust manifold).
Barrier: Lack of shielding between filter and the
Outcome
Actor E: Chief Engineer
E1, Chief Engineer, 15:15, Inspects generator temperatures and filters.
E2, Chief Engineer, 16:00, Noted high cooling water temperature alarm from port generator cylinder no.1.
E3, Chief Engineer, 16:00, Sounds general alarm.
E3, Chief Engineer, 16:00, Shuts down port generator.
Prior to the 20t July 1999
15:00
15:15
16:00

Figure 11.4: A MES Flowchart showing Causation in the Nanticoke Case Study

information that might otherwise be represented more directly by the events that stem from those conditions. He also argues conditions are often used to represent unsubstantiated factors that are difficult to validate after an incident has occurred. Others have argued that the MES approach is limited by the perceived complexity in developing and analysing the flowcharts [552]. As mentioned above, it can be difficult to identify a stable state for many complex technological systems. This, in turn, frustrates attempts to apply the P-Theory that drives MES analysis.

11.1.2 Sequentially Timed and Events Plotting (STEP)

The concerns mentioned at the end of the previous section led Hendrick and Benner to revise the MES approach [346]. Sequentially Timed and Events Plotting (STEP) provides a synthesis of ECF charting and MES [552]. It begins with the compilation of STEP cards. These provide an initial means of recording information about key events that occur during the course of an incident. They can be completed during any stage of a primary or secondary investigation. This reflects the concern that STEP should provide a pragmatic tool for investigators. It, therefore, attempts to avoid some of the notational excesses of the other analytical techniques that we have presented in previous chapters.

Event card identifier:	
Actor:	
Action:	
Time event began:	
Event duration:	
Data source/evidence:	
Event location:	
Description:	

Table 11.1: STEP card used to consolidate event information [346]

Table 11.1 illustrates the format of a STEP card. As can be seen, the information on these cards is closely modelled on the event blocks that support MES analysis. STEP cards do, however, record a number of additional items of information. In addition to the actor, time and action information that is captured by MES, STEP cards also introduce a free-text description of the event. They include information about the event location and its duration. Finally, STEP cards also record a Source identifier. This can be used to refer to the evidence that helped to identify the event. Such information can be useful when considering whether or not to support particular hypotheses about the course of events. The evidence that supports an event can be used to determine whether or not it should be retained in the face of competing explanations about the course of an incident.

Event information again provides the building blocks that are used to reconstruct the course of an incident. STEP relies upon a tabular format rather than the MES flowchart. The abscissa or vertical scale denotes the passage of time during an incident. The beginning and end of the accident sequence are, therefore, represented by the first and last columns in the matrix. Actors are represented on the ordinate, or horizontal, scale in the matrix rather than along the Y-axis in a MES chart. This tabular format offers a number of potential benefits during the initial stages of an investigation. Spreadsheets can be used to reduce the burdens associated with inserting new events and actors into an existing matrix. This might appear to be a trivial issue. As we have seen, however, the overheads involved in constructing graphical diagrams that involve many hundreds of nodes can dissuade investigators from using many of the more 'advanced' techniques that have been proposed to support incident analysis.

As mentioned before, STEP matrices do not include conditions. These were initially included to explain why an event occurred. Experience suggested, however, that investigators used conditions to introduce a range of biases into MES flowcharts. For instance, conditions were used to represent contextual factors that might not have had a direct impact upon the course of an incident. They can also be used to modify events so that they seem to be less significant that they might otherwise appear. The decision to exclude conditions from the STEP methodology was also justified by the observation that conditions are, typically, the result of previous actions. It can, therefore, be argued that they are superfluous to any subsequent investigation. Previous chapters have argued that conditions provide an important means of introducing some of the broader contextual factors that affect the course of many incidents. The decision to omit them from STEP matrices is, therefore, open to debate. It remains a continuing focus for on-going research into accident and investigation analysis. However, the following pages adopt the conventions introduced by the original STEP papers. Conditions are omited from the tabular representations of event sequences.

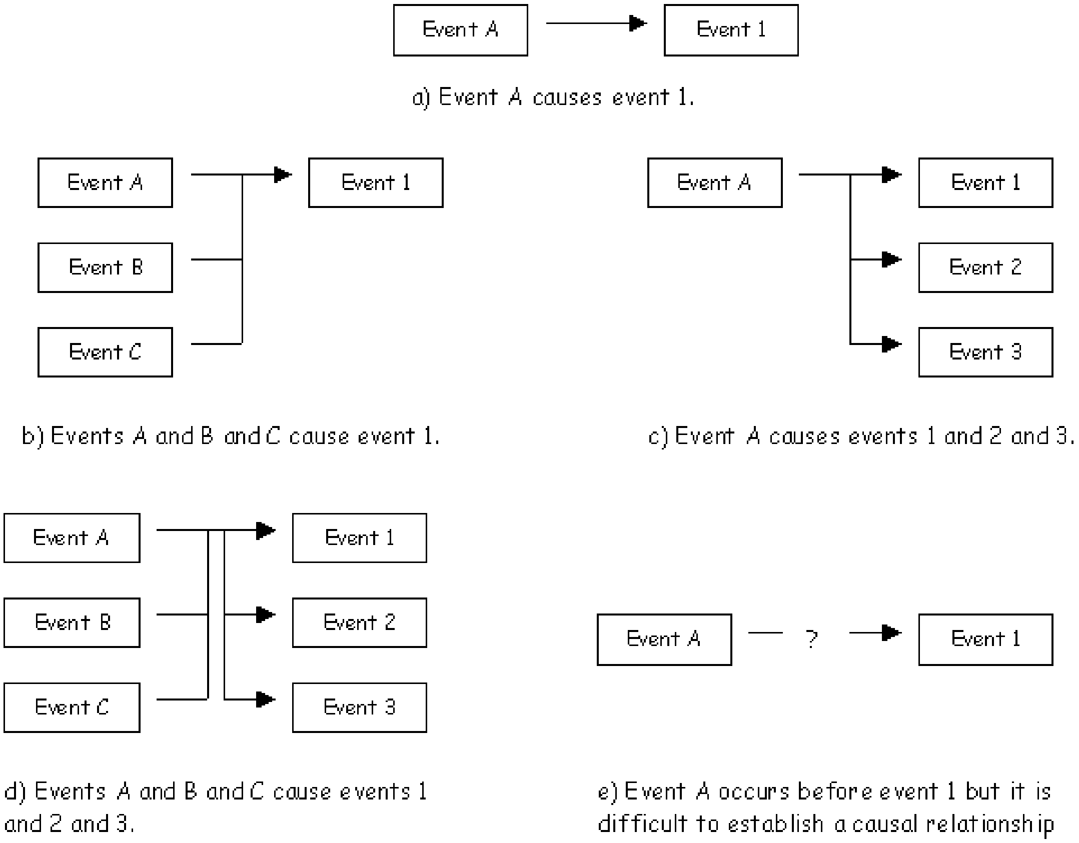

a) Event A causes event 1.

b) Events A and B and C cause event 1.

c) Event A causes events 1 and 2 and 3.

d) Events A and B and C cause events 1 and 2 and 3.

e) Event A occurs before event 1 but it is difficult to establish a causal relationship

Figure 11.5: Causal Relationships in STEP Matrices

The construction of a STEP matrix follows the P-Theory process outlined for MES analysis. The same consistency checks are also performed to ensure that the resulting worksheet provides a coherent temporal ordering over the events that it presents. The causal analysis of a STEP matrix also follows a procedure that is similar to that described for the MES flowchart. More recent expositions of STEP [74, 75] enumerate a broader range of causal relationships than appeared in the initial MES papers [72]. These are illustrated in Figure 11.5. Diagram a) denotes that event A is a direct cause of event 1. In other words, A is both a necessary and sufficient cause of 1. Diagram b) is similar to an AND gate within a fault tree. Events A, B and C are individually necessary for event 1 to occur. However, none of these events are sufficient for event 1 to occur unless A and B and C all occur at the times denoted by the abscissa. Diagram c) denotes a situation in which event A causes events 1, 2 and 3. This is important if the outcome of an event has an impact upon many other actors throughout a system. Events 1, 2 and 3 might represent these distributed, knock-on consequences. Diagram d) combines elements of diagrams b) and c) to denote that A, B and C are individually necessary and collectively sufficient for 1, 2 and 3 to occur. Finally, diagram e) denotes events that have a clear relationship in time but for which no causal explanation can be established. Such ambiguities form the focus for subsequent investigation of the underlying physical processes that characterise complex applications.

It is important to note that although diagram b) can be thought of as an AND gate, there is no equivalent of an OR gate within STEP matrices. If it is unclear what caused an event then investigators must introduce an event block that is labelled by a question mark. This is intended to avoid indicating "uncertainty about what happened" which is argued to be a weakness of the OR gate approach [75]. Whereas the use of events labelled by a question mark indicates "uncertainties in the description" [75]. It is difficult to interpret such distinctions. There are also pragmatic difficulties. Previous chapters have identified the limitations of current data recording devices. We have also described the problems associated with determining the causes of failure in hostile environments, such as space, where telemetry is strictly limited. This chapter does, however, follow the STEP conventions [75] . Disjunction are avoided.

Event card identifier: A1	
Actor:	?
Action:	Modifies forward filter cover/bolt sealing surface
Time event began:	Prior to 20th July 1999
Event duration:	?
Evidence:	Post incident inspection shows file marks are present on the cover/bolt sealing surface which was flat with no recess, unlike aft filter.
Event location:	Nanticoke forward fuel filter.
Description:	The copper washer gasket grooves are removed and this makes the sealing surface uneven.

Table 11.2: STEP card for the Nanticoke Filter Modification

Having introduced the underlying components of STEP, the following paragraphs apply this technique to analyse the Nanticoke case study. Matrices 11.2 and 11.3 present STEP cards that document information about key events. Investigators are intended to use these cards to help document the investigation progresses. Given the constraints of this case study, these cards were completed post hoc. They do, however, provide an illustration of the range of information that can be captured using these documents. For example, previous sections have explained the reasons why conditions are excluded from STEP matrices. This information can, however, be retained within the STEP card descriptions of key events. The condition labelled grooves for copper washer are removed, sealing surface is uneven and grooved with file marks in Figure 11.4 now forms part of the free-text description in Table 11.2.

STEP and MES are unusual in that they have been specifically intended to help investigators conduct a causal analysis during secondary, and even primary investigations. Other techniques, including ECF analysis and the application of Fault Trees, are far less explicit about when any causal analysis should begin. Many of the publications that propose the application of these approaches seem to make an implicit assumption that investigators have already secured any relevant information. We have argued in previous chapters that this is unrealistic. The identification of a potential causal factor can often lead to further investigation. For instance, if there is only circumstantial evidence that an event actually occurred. There is, therefore, a great deal to be learned from the comparatively simple documentary support offered by STEP cards. They avoid many of the maintenance overheads that are associated with the revision of more complex graphical and text-based analyses when new evidence becomes available.

The example STEP cards in Tables 11.2 and 11.3 illustrate further differences between this

Event card identifier: C2	
Actor:	Forward Fuel Filter
Action:	Starts spraying fine mist of fuel at pressure from the copper gasket.
Time event began:	After Chief Engineer's inspection at 15:15.
Event duration:	Until port generator shut-down at 16:00
Evidence:	When fire burns at high intensity, soot deposited on nearby surfaces is burnt off leaving a 'clean burn'. This is present slightly inboard of port generator valve covers 1 and 2; the general location of the fuel filters. Inspection of lubricating oil under the valve covers and two other starboard upper fuel filters rules out these sources.
Event location:	Nanticoke forward fuel filter.
Description:	If fuel was released under pressure from the copper gasket of the forward fuel filter then ignition could occur below the flash point of the fuel.

Table 11.3: STEP card for the Nanticoke Fuel Release

approach and the techniques that have been introduced in previous chapters. In particular, both include information about the evidence that supports the identification of particular events. The impact of modifications to the filter cover, event A1, is supported by a post incident inspection, which shows file marks are present on the cover/bolt sealing surface which was flat with no recess unlike the aft filter. The escape of fuel under pressure from the forward fuel filter, event C2, is supported by a more complex line of reasoning. When fire burns at high intensity, any soot that is deposited on nearby surfaces is burnt off leaving an area of 'clean burn'. Post incident inspections detected an area of clean burn slightly inboard of the port generator valve covers 1 and 2. This was in the general location of the fuel filters. These inspections also eliminated the possibility of the fire being fueled from three alternative sources. The importance of explicitly documenting such evidence should not be underestimated. The STEP approach encourages analysts to construct a single, 'deterministic' failure scenario. Disjunctions are not allowed when constructing STEP matrices from cards, such as those shown in Tables 11.2 and 11.3. Ambiguities are to be avoided as much as possible. Investigators must justify their analysis if their colleagues are to understand the evidence that supports the particular version of events that, in turn, supports any causal findings.

Figure 11.6 shows how a STEP matrix can be constructed to represent the causal relationships that exist between the various events that are described on STEP cards, such as those shown in Tables 11.2 and 11.3. As can be seen, there are strong similarities between this matrix and the MES flowcharts that were introduced in previous sections. However, there are no conditions. Some causal links have to be re-drawn because conditions are excluded from this form of analysis. For instance, in Figure 11.4 the modification event A1 led to a situation in which the sealing surface was uneven. This condition, in turn, affected the Watchkeeping Engineer's ability to obtain a fuel tight joint. In contrast, Figure 11.6 omits the condition. The modification event A1 might therefore have been shown as a direct causal link to event B2, which represents the Watchkeeping Engineer's attempts to obtain the fuel-tight seal. In contrast, Figure 11.6 shows that the modification event causes the fuel escape. This might seem like a subtle distinction but it reflects important differences between the MES and STEP techniques. In the former case, the initial event caused a condition that affected the Engineer's actions. Hence a causal link could be drawn from A1 to B2 through the mediating condition. In the STEP matrix, it cannot be argued that the modification event directly caused the Engineer to attempt to form a fuel-tight seal. Hence the modification event A1 and the Engineer's efforts, B2 contribute to the fuel release, event C2. The proponents of STEP argue that this clarifies

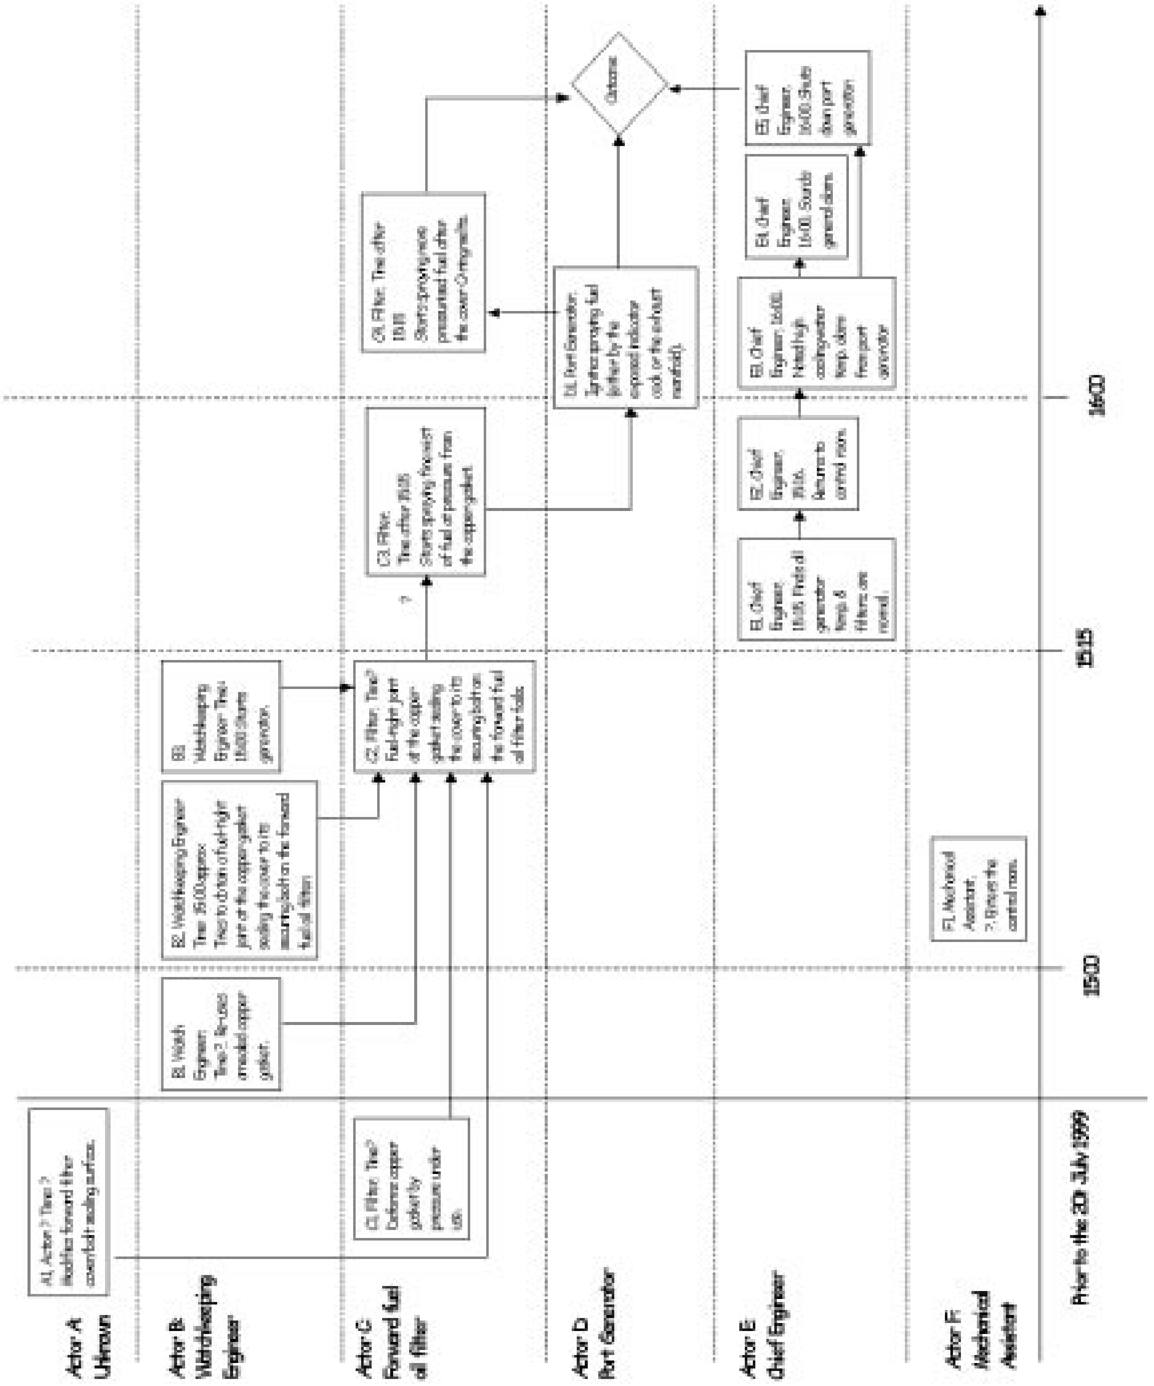
Actor A: Unknown
Actor B: Watchkeeping Engineer
Actor C: Forward fuel oil filter
Actor D: Port Generator
Actor E: Chief Engineer
Actor F: Mechanical Assistant
A1. Actor ? Time ? Modifies forward filter cover/bolt sealing surface.
B1. Watch Engineer. Time ?. Re-uses annealed copper gasket.
B2. Watchkeeping Engineer Time 15:00 approx Tries to obtain a fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter.
B3. Watchkeeping Engineer Time 15:00 Starts generator.
C1. Filter. Time? Deforms copper gasket by pressure under use.
C2. Filter. Time? Fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter fails.
?
C3. Filter. Time after 15:15 Starts spraying fine mist of fuel at pressure from the copper gasket.
C4. Filter. Time after 15:15 Starts spraying more pressurised fuel after the cover O-ring melts.
D1. Port Generator. Ignites spraying fuel (either by the exposed indicator cock or the exhaust manifold).
Outcome
E1. Chief Engineer, 15:15 Finds oil generator temp. & filters are normal.
E2. Chief Engineer, 15:16. Returns to control room.
E3. Chief Engineer, 16:00. Noted high cooling water temp. alarm from port generator
E4. Chief Engineer. 16:00. Sounds general alarm.
E5. Chief Engineer. 16:00 Shuts down port generator.
F1. Mechanical Assistant, ?. Enters the control room.
Prior to the 20th July 1999
15:00
15:15
16:00

Figure 11.6: STEP Matrix for the Nanticoke Case Study

the causal relationships between events. The condition into the MES diagram introduces a form of indirection between the modification event and the eventual fuel release that is not apparent in the STEP matrix of Figure 11.6. This practical standpoint is support by the more philosophical work of Lipton who argues that only events can be causes [496].

Figure 11.6 extends the previous MES analysis by considering a number of additional causal factors. In particular, the role of the Chief Engineer and the Mechanical Assistant are considered in greater detail. Events are introduced to denote that the Mechanical Assistant Enters the control room and that the Chief Engineer returns to the control room. These are used to explain why the fire was not detected until 16:00. This again illustrates how the application of causal analysis techniques continues to depend on the skill and expertise of the investigator. There is no automatic means of determining that these additional events ought to be introduced into a MES flowchart or STEP matrix. Table 11.4 illustrates how such insights may force investigators to develop additional STEP cards to represent information about a wider range of events. In this case, the card is used to record details about the Chief Engineer's return to the control room after his inspection at 15:15.

Event card identifier: E2	
Actor:	Chief Engineer
Action:	Returns to control room.
Time event began:	Approximately 15:16.
Event duration: Until high cooling water temperature alarm around 16:00.	
Evidence:	Witness evidence (Watchkeeping Engineer, Mechanical Assistant and Chief Engineer).
Event location: Engine control room.	
Description:	The Chief Engineer returned to the control room after observing that the generators and filters appeared to be functioning normally. The significance of this event is that they could not observe the port-side of the engine room from the control room. Neither the chief engineer nor the mechanical assistant made a visual inspection between 15:15 and 16:00 and this gave the fire an opportunity to take hold.

Table 11.4: STEP card for the Chief Engineer's Monitoring Activities

Table 11.4 also illustrates a number of problems that complicate the application of the STEP approach. Firstly, the card explains the significance of the Chief Engineer's decision to return to the control room. He could not observe the port-side of the engine room from the control room and, therefore, was unlikely to directly observe the fire until it had taken hold. This information is not included on the STEP matrix in Figure 11.6. This introduces cross-referencing problems that affect the use of multiple representations for the same events. Investigators must not only understand the causal relationships represented on the matrix but they must also follow the more detailed information that is represented on each of the cards. This might seem to be a relatively trivial demand. It can, however, impose significant burdens when STEP matrices are used to represent complex, safety-critical incidents involving several hundred events.

There are further problems. The STEP card in Table 11.4 records that neither the chief engineer nor the mechanical assistant made a visual inspection between 15:15 and 16:00 and this gave the fire an opportunity to take hold. The previous STEP matrix does not document this temporal information. One solution would be to introduce an additional field into a STEP card. This would distinguish the duration of an event and from the duration of its effects. For example, the Chief Engineer only took a few seconds to enter the control room but they remained there until 16:00. Such additions to the STEP card introduce further problems. Events often trigger a number of different effects. The ignition event created heat and smoke, it also eventually triggered temperature alarms. Each of these effects have different durations. The smoke and heat were eventually countered by ventilating the engine room after Halon gas had been used to extinguish the fire. The alarms continued until the ship had been secured. The tractability of the STEP cards approach would clearly be sacrificed if investigators had to introduce this duration information into the more concise summaries of key events.

These overheads can be avoided by explicitly introducing stopping events into a STEP matrix. The continued presence of the Mechanical Assistant and the Chief Engineer in the control room can, therefore, be inferred by the absence of any event to denote that they left the control room. Such inferences carry a degree of uncertainty. Investigators may forget to introduce these terminating events. Figure 11.6 uses event C4 to denote that pressurised fuel begins to spray at an increased rate when the filter cover O-rings melt after 15:15. We have not, however, specified when this fuel release ended. In constructing the STEP matrix it was assumed that investigators would recognise that the release ended when the engineer shut-down the port engine at 16:00. Unless explicit stop-events are introduced, there is a danger that investigators may rely upon incorrect inferences about the duration of key properties during an incident. These problems should not be surprising. The difficulty of representing events and duration also affected the time-lines introduced in Chapter 9 and the ordinate scale of the STEP matrix can be viewed as a time-line.

Previous sections have explained how both MES and STEP derive directed graphs of an incident. Nodes represent events in STEP, or events and conditions in MES. Edges represent the causal relations that hold between nodes in the graph. We have not, however, described how investigators can identify root causes from the various causal factors that are used to construct these graphs. One solution would be to replicate the analytical techniques that were introduced a the end of Chapter 10. In addition to the validation steps, which ensure that causal factors are both necessary and sufficient, analysts must distinguish those events that represent more general (root) causes from those that characterise a particular incident. It is important to note, however, that the developers of the STEP and MES techniques have been highly critical of previous attempts to derive methods for root cause analysis. Benner, in particular, has argued that attempt to distinguish root causes from causal factors can misdirect investigators to find a few ‘silver bullets’ instead of understanding, describing and explaining the entire incident process [76]. He goes on to argue that root causes are often ‘judgemental, unverifiable conclusions’ that typically cannot be validated by ‘objective quality controls’. These comments are consistent with the STEP focus on determining the particular events that contributed to an incident. Conditions that might represent wider causal factors are deliberately excluded from this approach. In contrast, STEP focuses on the evidence that supports the introduction of particular events into the associated matrices.

P-Theory suggests that investigators must focus on the initial perturbation that causes any subsequent failure. Figure 11.6 starts with the initial modifications to the forward filter cover/bolt sealing surface. P-Theory also suggests that investigators should consider causal events that compromise protective barriers. For example, the Watchkeeping Engineer might have reported the problems experienced in fitting the gasket. These events represent missed opportunities for the system to return to an initial ‘homeostasis’. This focus on the particular causes of an incident provides important benefits. It is intended to reduce the likelihood that external pressures will ‘persuade’ investigators to introduce arbitrary contextual factors, or conditions, as a means of explaining particular events [74]. There is, however, a danger that the application of MES and STEP will miss important underlying causes of an incident. For example, previous sections have argued that organisational and managerial failures can jeopardise a number of different barriers. These failures can be analysed and measured, for instance in terms of participation rates in incident reporting schemes or in the number of regulatory sanctions that were previously applied to an organisation. It is unclear how such factors might be represented as causal events within a STEP analysis.

A number of further problems complicate the application of STEP [74, 75]. These include limitations that affect this particular approach. They also include more general issues that affect all forms of causal analysis:

1. *Incomplete chains between the first and last events.* If it is not possible to establish a path through the causal connections in the matrix then investigators must seek additional evidence about events that might not have been identified. This can involve the use of additional techniques, such as Fault Trees or the Change analysis and Barrier Analysis methods that were used in conjunction with ECF charts. Alternatively, the scope of any report might be confined to those events that can be accurately identified from the available evidence. This clearly jeopardises the insights that might have been obtained from any analysis of the incident. Investigators must, typically, take steps to increase the amount of 'diagnostic' information that can be obtained from any potential future incidents.

2. *Unconnected events after the causal analysis.* The developers of the STEP method argue that investigators must avoid unconnected events. For example, events F1 in Figure 11.6 denotes that the Mechanical Assistant enters the control room. It does not, however, have any direct causal relationship with the subsequent events in the Nanticoke case study. It has been argued that, at best, these unconnected events can divert investigator's attention away from more important causal sequences. Scare development resources can be allocated to deal with these extraneous peturbations that need not have affected the course of an incident. At worst, it is argued that they provide "handles for others to grasp to raise irrelevant, unnecessary and invalid questions about the accident" [75]. It is argued that investigators should delete these unconnected events from a STEP matrix because they can mislead rather than enlighten other investigators. The dangers with such a policy are clear. Investigators run the risk of deleting information that might enable their colleagues or other readers to identify important causal relations that might have been overlooked in any previous analysis. If analysts follow this advice then there ought to be some documentary evidence to record their decision so that others can follow the justification for removing information from the matrix.

3. *Inconsistent data requirements.* The increasing inter-connection and functional sophistication of safety-critical systems poses considerable challenges for incident analysis. This complexity has been exacerbated by the increasing recognition that more and more factors ought to be considered during any investigation. The scope of any analysis has broadened beyond individual operator error and component failure to examine more systemic causes of incidents and accidents. It is not surprising, therefore, that analytical techniques such as STEP should yield complex accounts of the mishaps they represent. This can lead to conflict if managers expect 'simple' descriptions of complex failures. Further problems can arise if the products of a STEP analysis do not correspond to the categories expected by a regulators reporting system. As Benner notes, "this very frequent problem often arises after statisticians design forms for data collection, then declare that the statistical elements on the forms are significant investigative data and train investigators to 'fill out the form' rather than investigate the accident" [75]. Later sections will assess these problems in greater detail. For now it is sufficient to recognise that they stem from the wider organisational and regulatory environment that surrounds an incident reporting system. Investigators must clearly be aware of such issues before attempting to pioneer the introduction of new analytical techniques.

This section has identified that changes that have been introduced between earlier version of the MES analysis technique and the more recent STEP approach. MES and STEP can be seen as variants of the same underlying ideas. Both rely upon the notion of event blocks that are associated with actors and can be mapped onto a time-line. These similarities should not be surprising given that STEP extends Benner's earlier work on MES [72]. Some confusion has arisen because these two different terms have been used synonymously. Investigators have referred to MES when applying the tabular forms associated with the techniques in the STEP handbook [346]. We have attempted to make a clear distinction between these techniques, however, readers should be aware of the potential confusion given these strong similarities.

11.2 Check-List Approaches

Previous sections have focussed on event-based techniques that encourage analysts to reconstruct or model the development of an incident over time. A number of alternative techniques have, however, rejected this approach. In contrast, they often assume that analysts develop and maintain more implicit models of the events that contribute to an incident. This arguably reflects a more pragmatic attitude to the partial nature of the evidence that is available in the aftermath of many mishaps. These approaches instead provide checklists that prompt investigators to look for a number of predefined features that are common to a wide range of incidents and accidents. The US National Patient Safety Foundation's (NPSF) report on the 'Scientific Basis for Progress on Patient safety' summarised the strengths and weaknesses of these approaches:

> "Collections of incidents and accidents cry out for classification. The apparent similarities and differences between the events, their outcomes, and the circumstances that precede them encourage us to organise them in categories and rank them in severity. But classification also has its own hazards, especially in complex domains where there are multiple possible paths to any outcome and multiple possible outcomes from any path. Classification involves identifying relevant similarities and differences; their effective use depends on being able to know a priori what relevant means... Classification does involve a type of analysis but a type that greatly constrains the insights that can be obtained from the data. Typically, when classification systems are used as the analysis, a report of an incident is assigned, through a procedure or set of criteria, into one or another fixed category. The category set is thought to capture or exhaust all of the relevant aspects of failures. Once the report is classified the narrative is lost or downplayed. Instead, tabulations are built up and put into statistical comparisons. Put simply, once assigned to a single category, one event is precisely, and indistinguishably like all the others in that category." [181]

The following paragraphs use a number of different causal analysis techniques to illustrate and expand on these observations. In contrast, Chapter 15 describes how checklist approaches to causal classification can also be used as the indices in information retrieval systems.

11.2.1 Management Oversight and Risk Tree (MORT)

Figure 11.7 illustrates the Management Oversight and Risk Tree. This is the central component of what has become known as MORT [429]. As can be seen, MORT diagram is constructed using the elements of a fault tree. An undesired event can be either the result of oversights and omissions or it is the result of an assumed risk. Assumed risks "are defined as only those risks that have been analysed and accepted by the proper level of management; unanalysed or unknown risks are not considered to be Assumed Risks" [203]. If an oversight or omission has occurred then it can be categorised as being the result of either a management failure or of a failure in specific technical controls. If there has been a break-down in management then either there was a failure in the implementation of some policy or the policy was flawed or the risk assessment was less than adequate. A failure in management risk assessment can occur if incorrect goals were established for a project or the information systems used to support a risk assessment were less than adequate or the hazard analysis process was flawed or the safety review program was less than adequate. As can be seen, the components of the MORT diagram provide a check-list that can be used to analyse and categorise the potential causes of an incident.

The MORT diagram was intended to provide a template that might guide the causal analysis of incidents and accidents. There is an obvious danger that investigators will force an incident to fit one or more of the categories in the MORT checklist. The proponents of this approach have responded by extending the range of factors that are included in the MORT diagram. For instance, the version of Figure 11.7 includes over 1,500 basic events. This leads to a difficult trade-off. By extending the scope of the MORT diagram, investigators are more likely to find an appropriate causal factors that describes their incident. By extending the scope of the MORT diagram, investigators may

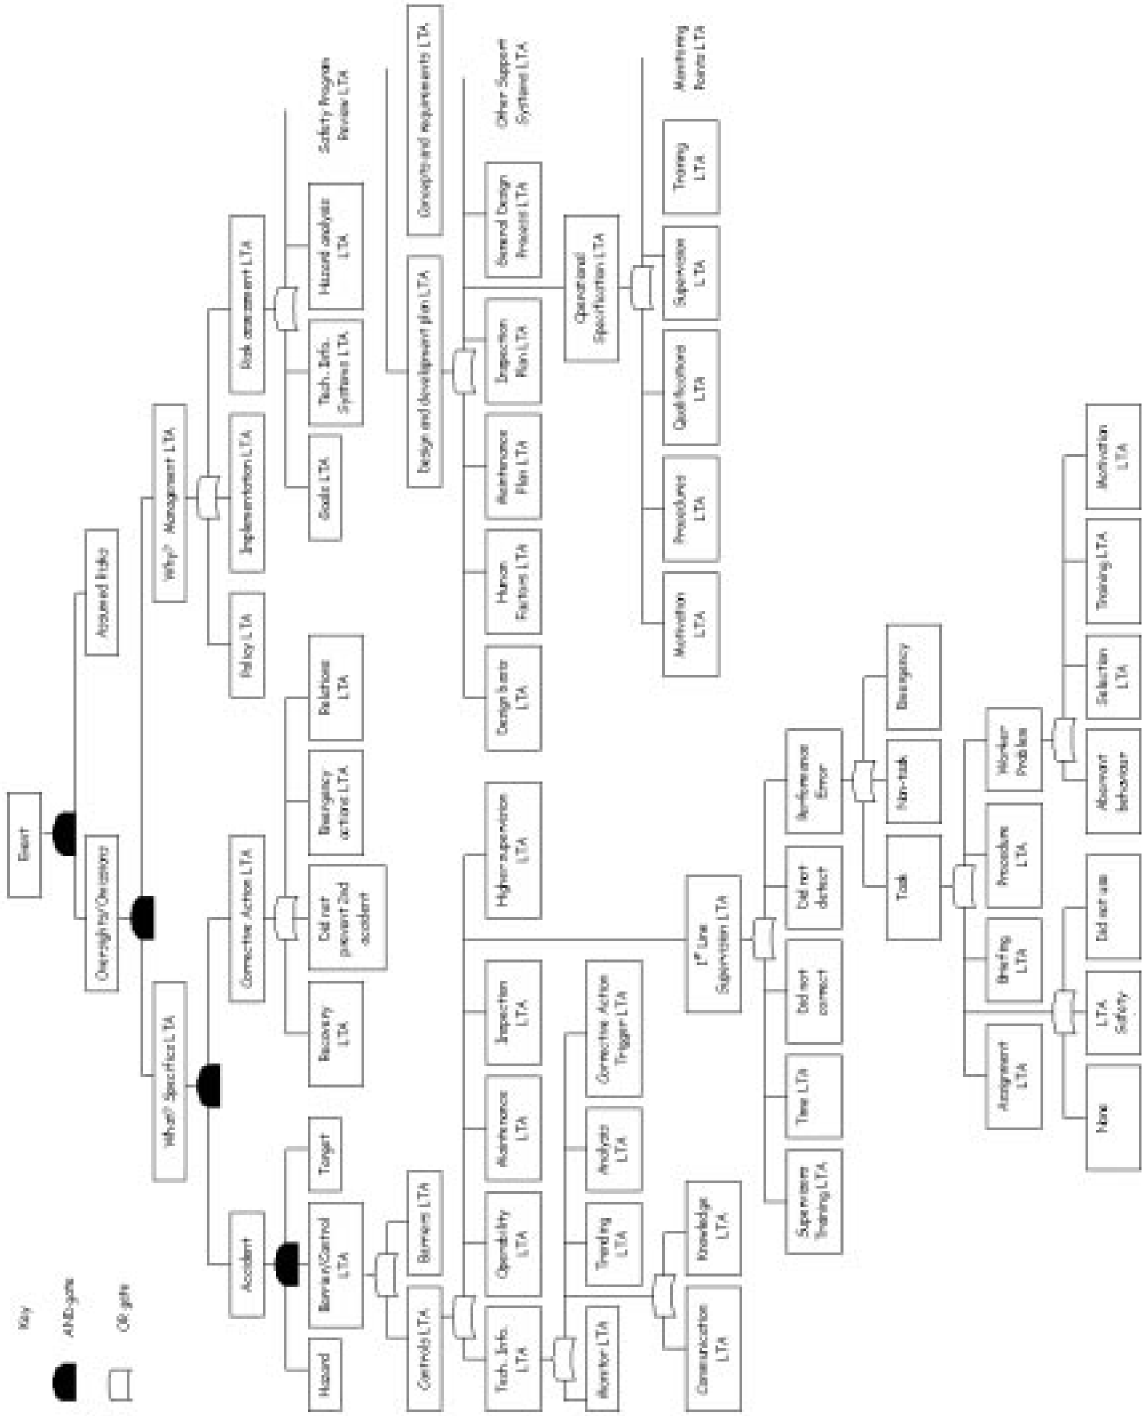

Figure 11.7: The Mini-MORT Diagram

also experience more difficulty in distinguishing between the many different forms of failure that are described by each of the leaf nodes. In consequence, the US Department of Energy advocates the use of a stripped-down version of the full MORT diagram [204]. This mini-MORT provides approximately fifty basic events but each of these denotes a far broader set of causal factors than the more detailed versions of the diagram.

MORT diagrams embody their developers' view of accident causation. The branches of the tree reflect a concern to assess management responsibility. There is also provision for assessing the technical context in which an incident occurs. Human factors issues are also captured, arguably in a rather narrow fashion, by focusing on errors of commission . A further branch traces the failure of barriers. As a result, the barrier analysis introduced in Chapter 10 is often used as a precursor to MORT classification. There is also a preoccupation with understanding and assessing the causes of any potential energy release [298]. One consequence of this is that MORT also provides an implicit definition of incidents and accidents. An accident occurs if an unwanted energy flow affects a vulnerable target. An incident occurs if an unwanted energy flow occurs without hitting such a target [457]. This is consistent with the use of barrier analysis and reflects their common origin within the nuclear industry. Johnson developed most of the MORT approach while working for the US National Safety Council and under a contract from the US Atomic Energy Commission [429]. As mentioned, the US Department of Energy continues to advocate this method [204, 203]. The MORT approach, therefore, combines concepts from management and from safety analysis. It captures the notion that management has a profound impact upon the effectiveness of barriers that prevent unplanned energy releases.

MORT analysis consists of two principle stages. Firstly, analysts must consider what happened during an incident. This involves a traversal of the what? sub-tree under the oversights and omissions branch. This is intended to help the analyst identify the barrier or control problems that contributed to the incident. Secondly, the analyst must then identify any management elements on the why branch of the MORT diagram that contributed to these particular problems. It is important to document each of the problems and summarise the findings of the analysis.

This process of iteratively describing what happened and then searching for causal explanations in the why branch is guided by a number of questions that analysts can ask as they inspect each node in the MORT diagram. For example, the US Department of Energy MORT user guide provides the following question that can be asked to determine whether or not any emergency response was adequate. This corresponds to the leaf node with the following pathEvent: Oversights and Omissions: What? : Corrective Actions: Emergency Actions:

> "*Emergency Action (Fire Fighting, Etc.) Less Than Adequate* Was the emergency response prompt and adequate? Which emergency response teams were required? Were they notified and did they respond? [Include local facility fire brigade, health physics team, fire department, bomb squad, and other speciality teams. Be sure to consider delays or problems in both notification and response.] " [203]

These questions appear, at first sight, to be relatively straightforward. Unfortunately, a number of factors complicate this analysis. The use of the term 'less than adequate' implies a value judgement. There can often be considerable disagreement about what does, and what does not, represent an adequate response. Even in countries that publish national guidelines for response times, there can be considerable debate about whether the nature of any response was appropriate given the scale of an incident [217, 210]. Some investigators, including Benner [72], argue that these value judgements are open to political pressure and bias in the aftermath of safety-critical incident.

Even with the additional complications created by the validation of value judgements, the previous question is relatively simple in contrast to some of the other guidelines that are intended to support MORT analysis. This point is illustrated by the following questions. These are intended to guide the analysis of a supervisor's failure to correct a hazard. Each of these questions relates to further basic events that are present in more complete versions of the MORT diagram. They would be shown under Event: Oversights and Omissions: What? : Accident : Barrier/Controls/ Controls/1st Line Supervisor/ Did Not Correct Hazards in Figure 11.7:

> *Did Not Correct Hazards:* Was an effort made to correct the detected hazard?

- Interdepartment Coordination Less Than Adequate: If the accident/incident involved two or more departments, was there sufficient and unambiguous coordination of interdepartment activities? [Interdepartment coordination is a key responsibility of the first line supervisor. It should not be left to work level personnel.]
- Delayed: Was the decision to delay correction of the hazard assumed by the supervisor on behalf of management? Was the level of risk one the supervisor had authority to assume? Was there precedent for the supervisor assuming this level of risk (as then understood by him)? [Note a decision to delay correction of the hazard may or may not transfer to the Assumed Risk branch. It was an assumed risk only if it was a specific named event, analysed, calculated where possible, evaluated, and subsequently accepted by the supervisor who was properly exercising management-delegated, decision-making authority.]
- Was the decision to delay hazard correction made on the basis of limited authority to stop the process?" [203]

The previous two examples have illustrated the questions that can be used to guide the analysis of the what sub-branch in a MORT analysis. Previous paragraphs have, however, argued that investigators must also identify the reasons why these events occurred. This involves an analysis of the why sub-branch under the oversights and omissions node. Questions can again guide this form of analysis. For example, the following guidelines corresponds to the leaf node with the following path Event: Oversights and Omissions: Why? : Management : Risk Assessment : Safety Program Review : Design and Development Plan : Human Factors. They direct an analyst to consider the impact that a managerial failure to consider human factors issues may have had upon the course of an incident:

> "*Human Factors Review Less Than Adequate:* Has consideration been given in design, plan, and procedures to human characteristics as they compete and interface with machine and environmental characteristics?
>
> - Professional Skills Less Than Adequate: Is the minimum level of human factors capability, needed for evaluation of an operation, available and will it be used? (275)
> - Did Not Describe Tasks: For each step of a task, is the operator told: When to act? What to do? When the step is finished? What to do next? (276)
> - Allocation Man-Machine Tasks Less Than Adequate: Has a determination been made (and applied) of tasks that humans excel in versus those tasks at which machines excel?
> - Did Not Establish Man-Task Requirements: Does the review determine special characteristics or capabilities required of operators and machines?
>   - Did Not Define Users: Is available knowledge about would-be users defined and incorporated in design?
>   - Use of Stereotypes Less Than Adequate: Are checklists of stereotypes (typical, normal, expected behaviour) used in design? (e.g., Is a control turned right to move a device to the right?) Are controls coded by size, colour, or shape?
>   - Displays Less Than Adequate: Are displays used which can be interpreted in short time with high reliability?
>   - Mediation Less Than Adequate: Is consideration given to delays and reliability of interpretation/action cycles?
>   - Controls Less Than Adequate: Are controls used which can be operated in short times with high reliability?
> - Did Not Predict Errors: Is there an attempt made to predict all the ways and frequencies with which human errors may occur, and thereby determine corrective action to reduce the overall error rate?

- Incorrect Act: Have all the potential incorrect acts associated with a task been considered and appropriate changes made?
- Act Out of Sequence: Has the consequence of performing steps of a task in the wrong order been considered and has appropriate corrective measures been made?
- Failure to Act: Is there an attempt to reduce the likelihood of operators omitting steps or acts which are required by procedure?
- Act Not Required: Are all the steps that are needed to accomplish a task required in the procedures? Are only those steps in the procedure?
- Malevolence: Are deliberate errors and other acts of malevolence anticipated and steps taken to prevent them or reduce their effect?" [203]

The MORT user guidelines emphasise a number of additional practical observations that have emerged from the application of this technique during incident and accident investigations [203]. The approach works best if it is used to focus discussion and debate. Any figures or forms that are produced during the analysis should be considered as working documents that can be revised and amended as work progresses. MORT, therefore, provides analytical guidance; 'it helps avoid personal hobbies, bias, or the tunnel vision that commonly results from pet theories of accident causation' [203]. It should not be seen as a framework to be imposed upon a final report. It can, however, be used as a quality control mechanism to identify any potential omissions in a final report. Investigators can use the questions to ensure that they have described both what happened and why those events occurred. Finally, experience in applying MORT has shown that even the full version of the diagram cannot cover all aspects of some incidents. If a mishap is not covered in any of the branches then analysts are encouraged to extend the existing diagram using the basic fault tree gates that were introduced in Chapter 9.

Having raised these caveats, it is possible to illustrate the application of MORT to the Nanticoke case study that was introduced in previous sections. As mentioned above, MORT analysis begins by determining what happened during an incident or accident. Investigators traverse the what branch of the tree, such as that shown in Figure 11.7, asking whether or not each potential failure contributed to the incident under investigation. MORT assumes that investigators have sufficient evidence to perform such an analysis. It does not provide any explicit guidance on how to go about satisfying this prerequisite, however, others have extended the approach to provide this support [444]. At the highest level, this traversal of the MORT diagram encourages investigators to identify the hazard that threatened potential targets within the system [298]. In our case study, the hazard can be identified as the danger of a fire being started by a pressurised fuel release from a fuel filter onto the adjacent indicator tap or uncovered exhaust manifold. This hazard threatened a number of different targets. Most immediately it posed a danger to the people and systems in the engine room. Ultimately, it threatened everyone on the vessel and even other ships that were operating in the same area as the Nanticoke.

As can be seen from the left sub-branches of Figure 11.7, analysts must also identify the ways in which any barriers or controls were circumvented during an incident. Barriers typically protect or shield a target from a hazard. Controls make it less likely that a hazard will occur in the first place. These terms are, however, often used interchangeably [552]. This imprecision is justified by the practical difficulties of distinguishing between these two different forms of defence. For instance, more regular inspections of the filter assembly might have made the fire less likely. Crew members might have noticed the leak before ignition. More frequent inspections might also have acted as a barrier by raising the alarm as soon as the fire had started. The practical problems of distinguishing between these different forms of protection helps to explain an imbalance in the MINI-Mort tree of Figure 11.7. This diagram provides considerable detail about the potential forms of control failure. This level of detail is not, however, reflected by the portion of the tree that considers inadequate barriers. This imbalance is also justified by the observation that these failures often take similar forms. Inadequate technical information or maintenance procedures can threaten both of these potential defences.

Barriers prevent hazards from having adverse consequence once they occur. They can be thought of as protection devices or shields that guard the target from the hazard. It can be argued that the barriers worked well in the Nanticoke case study because the fire was ultimately extinguished without loss of life or serious injury. Conversely, it can be argued that the barriers failed because the ship suffered considerable damage. The relatively limited fire managed to burn through the common cable tray that contained all of the steering systems. After 1st September 1984, duplicated steering power and control systems had to be routed as widely as possible throughout a vessel so that an isolated fire was unlikely to destroy all of these redundant systems. The Nanticoke was built in 1980 and so lacked the protection offered by the 1984 requirement. In consequence, the vessel was effectively disabled until an alternative power source could be rigged to the steering gear.

As mentioned, controls make it less likely that a hazard will occur. Figure 11.7 documents a number of potential weaknesses that can jeopardise adequate control. For example, the Nanticoke incident was arguably caused by inadequate maintenance. The modifications to the forward filter cover and bolt sealing surface left grooves that made it hard for the watchkeeping engineer to achieve a fuel-tight joint. This analysis shows how the MORT diagram can be used as a check-list to guide the analysis of what happened during an incident. It also illustrates some of the complexity that frustrates the use of checklist techniques. Damage to the seating surface not only suggests inadequate maintenance, it also indicates that there may have been inspection problems. Crew members might have recognised the potential for a fuel leak during previous rounds of preventive maintenance. This illustrates the way in incidents often stem from the failure of several different controls. Problems arise if investigators form different opinions about the salience of these failures. For instance, some analysts might discount the importance of inspection failures by arguing that the true significance of the seating damage could only have been determined with hindsight. Other analysts might stress the importance of inspection failures by arguing that the watchkeeping engineer should have reported their problems in obtaining a fuel tight seal during the maintenance that immediate preceded the incident. Such differences of interpretation make it very important that analysts both document and justify the findings of their MORT analysis. These justifications can then be reviewed and challenged before any subsequent causal analysis.

There are a number of differences that distinguish checklist approaches, such as MORT, from event-based techniques, such as STEP and MES. In particular, checklist approaches often abstract away from the temporal properties that are a central concern of the flowcharts and tabular forms in previous sections. The initial stages of a MORT analysis identify instances of generic failure types. They do not chart the timing of events. This is both a strength and a weakness. The MORT diagram cannot, in isolation, be used to reconstruct the way in which an incident developed over time. There is, therefore, no guarantee that investigators will identify omissions or inconsistencies in the events leading up to an incident. On the other hand, previous sections have criticised event-based techniques that force analysts to model precise event sequences which are unlikely to recur in future incidents. The identification of MORT failure types can, in contrast, generalise from the particular observations that characterise an individual incident. There are further benefits. By abstracting away from temporal properties, the MORT classification process can help investigators to identify similarities between latent and catalytic failures. Such similarities can be difficult to demonstrate with event-based techniques that deliberately separate the presentation of events that occur at different times during an incident. For instance, inadequate inspections may have contributed to the latent conditions behind the Nanticoke incident. Crew members failed to recognise the damage to the seating surface and this ultimately made it difficult for the engineer to achieve a fuel-tight seal. Inspection failures also characterised more immediate events during the incident. The engine room was not inspected between 15:15 and 16:00. Subsequent analysis might determine that these different failures had very different causes. The key point is, however, that the MORT style of analysis can help to identify potential similarities between failures that occur at different times during the same incident.

As with any checklist approach, MORT provides prompts that encourage analysts to consider a broad range of potential failures that might contribute to incidents and accidents. For example, the Nanticoke case study partly stemmed from operability problems. There were no new copper gaskets. Once a used copper gasket has been deformed by use, it is more difficult to obtain a tight seal for subsequent use even if it has been annealed. Other failures can be associated more directly with individual operators. For instance, the Mini-MORT diagram of Figure 11.7 includes a branch that represents inadequate intervention by the first line supervisor. As we have seen, it can be argued that they failed to correct the damage incurred during previous modifications to the filter. It can also be argued that they failed to detect the leak or the fire before it had taken hold.

Figure 11.7 also shows how further branches focus on the response to an incident. For instance, it can be argued that the emergency actions that were taken in response to the incident were complicated by the lack of any emergency exit from the control room. In consequence, the chief engineer had to follow hand rails out of the engine room. The corrective actions branch of the MORT diagram also includes a node Did not prevent 2nd accident. This supports the analysis of incidents in which the same hazard occurs more than once. For example, the fuel might have reignited after the initial fire had been extinguished. More widely, this node can encourage investigators to consider whether an incident forms part of a wider pattern. Chapter 15 will stress the importance of such activities. Investigators must look beyond the immediate response to an incident in order to learn from previous attempts to address similar failures. For example, the Transportation Safety Board of Canada identified that four similar engine room fires had occurred on Canadian ships within six months of the Nanticoke incident [621]. Previous ship safety bulletins had not resulted in adequate barriers being placed between potential fuel sources and adjacent exposed, hot surfaces. Subsequent analysis of the reasons why the fire occurred must, therefore, explain this failure to act upon previous safety bulletins.

Sub-Tree: What/Accident

What?	Rationale
Hazard	Danger of a fire being started by a pressurised fuel release from a fuel filter onto the adjacent indicator tap or uncovered exhaust manifold.
Targets	People and systems in the engine room. Everyone on the vessel. Other ships in the same area as the Nanticoke.
Barriers	
Did not use	More frequent inspections might raised the alarm sooner.
Did not provide	Fire burnt through common cable tray containing all of the steering systems. Nanticoke was disabled until alternative power source was rigged for steering gear.

Table 11.5: MORT (Stage 1) Summary Form for Hazard, Targets and Barriers

Tables 11.5 and 11.6 summarise the results of the first stage in the MORT analysis of the Nanticoke case study. These tables are intended to provide a focus for discussion. Previous paragraphs have argued that considerable disagreements are possible over our interpretation of which nodes best capture the failures that contributed to this incident. It is also important to notice that Table 11.5 extends the Barrier branch from Figure 11.7. The nodes Did not use and Did not provide reflect types of failure that were described as part of the introduction to barrier analysis in Chapter 10. This illustrates the way in which analysts may have to extend the pre-defined categories within a Mini-MORT diagram. In this case, however, these additional nodes are consistent with those included in the full MORT diagram.

Previous sections have described how this first stage of identifying *what* happened helps to drive a more detailed causal analysis of *why* those failures occurred. Before making this transition, however, it is possible to make a few observations about the use of MORT to drive an initial assessment of the Nanticoke case study. As we have seen, there is no automatic or semi-automatic procedure for

Sub-Tree: What/Accident

What?	Rationale
Controls	
Inspection LTA	More regular inspections of filter assembly might reduced likelihood of fire. Crew members might have noticed the leak before ignition.
	Crew members (arguably) might have reported problems in obtaining a fuel tight seal during maintenance immediately before the incident.
	Engine room was not inspected between 15:15 and 16:00.
Maintenance LTA	Modifications to forward filter cover and bolt sealing surface left grooves that made it hard to achieve a fuel-tight joint.
Operability problems	No new copper gaskets. Copper gaskets are deformed by use and pose more problems in obtaining a tight seal even if they have been annealed.
1st Line Supervision LTA	Failure to identify and correct damage incurred during previous modifications to the filter.
	Failed to monitor engines during interval prior to the fire (15:15 to 16:00).
Emergency actions LTA	No emergency exit from the control room. Chief engineer had to follow hand rails out of the engine room.
Did not prevent 2nd accident	Four similar engine room fires occurred on Canadian ships within six months of the Nanticoke incident. Ship safety bulletin (13/85) had not resulted in adequate barriers being placed between potential fuel sources and adjacent exposed, hot surfaces.

Table 11.6: MORT (Stage 1) Summary Form for Controls

identifying the particular failures that characterise an incident or accident. In contrast, investigators must rely on subjective judgement and prior expertise to determine which of the MORT nodes most accurately describe what led to the incident. There are no guarantees that different investigators will derive similar classifications for the same incident. This would seem to be unlikely given that particular conditions, such as the damage to the seating, can be the result of several inadequacies throughout the left-hand branch of the MORT diagram. The proponents of this approach have argued, however, that MORT provides a focus for discussion rather than a method for deriving a definitive analysis or single interpretation of events. This is an important observation given that there can be considerable disagreement not simply about the course of an incident but also about the precise meaning of each category within the MORT diagram. As we have seen, investigators often experience considerable practical difficulties in distinguishing between a barrier and a control. Some organisations have responded to these potential problems by developing considerable in-house documentation to support the use of MORT [203]. This material includes training material, case studies and style guides that reflect a particular approach to the MORT technique. Others have gone further. For instance, Kjellén has extended MORT to develop SMORT (Safety Management and Organisation Review Technique) [444]. This provides explicit support for data collection during incident investigations. As we have seen, this support was not part of the initial MORT approach. Such elaborations combined with explicit encouragement to extend the MORT diagram if it does not capture key aspects of an incident have resulted in a situation in which the term MORT is often used to describe a very varied collection of subtly different techniques. These techniques vary both in the checklists that are used and in the supplementary methodological support that is provided to guide their application.

A number of further observations can be made about the Nanticoke case study. The MORT diagram illustrated in Figure 11.7 captures the emphasis that this technique places upon failure. The diagram prompts investigators to identify what went wrong by systematically considering the ways in which various aspects of performance were less than adequate. Previous chapters have, however, argued that near-miss incidents often provide vital information about those barriers and controls that worked effectively to prevent an accident from occurring. For example, the Halon system on the Nanticoke provided an effective final resort after the crew made two unsuccessful attempts to fight the fire themselves. It can, therefore, be argued that investigators ought to repeat their analysis of a MORT diagram to identify these mitigating factors whose performance was At or Beyond Expectation (ABE) and not Less Than Adequate (LTA).

The second stage of MORT analysis helps investigators to determine the causes of an incident. This is done by identifying those elements in the why branch that contributed to each of the failures that were summarised in Tables 11.5 and 11.6. At the highest level, the overall hazard was the danger of a fire started by a pressurised fuel release from a fuel filter onto the adjacent indicator tap or uncovered exhaust manifold. It can be argued that this was the result of an inadequate risk assessment. The operators and crewmember failed to recognise the potential threat to everyone on the vessel and to other ships in the area. As before, the MORT diagram can be used to guide the analysis of what might have caused this failure. The Risk Assessment LTA branch contains a number of detailed nodes that investigators can adopt as working hypotheses about the factors that led to an incident. For example, Table 11.5 argued that more regular inspections might have prevented the fire from developing if the crew had been able to raise the alarm sooner than they did. The failure to effectively implement such a barrier can be explained in terms of the node Inspection Plan LTA which is located under the path Why? Management LTA : Risk Assessment LTA : Safety Program Review LTA : Design and Development plan LTA in Figure 11.7. Similarly, the failure to provide a sufficient barrier to protect the control cables for the steering system can be explained in terms of the Design basis LTA node which appears at the same level as Inspection Plan LTA. Had the Nanticoke been built after the September 1984 regulations were introduced then the cables would have been distributed more widely throughout the vessel. An isolated fire would then have been less likely to damage all of the redundant steering systems.

Investigators can also use the MORT diagram to identify potential reasons why Controls failed to protect the system. For example, Table 11.6 suggested that inspections might have been less than adequate because crewmembers might have noticed the possibility of a leak well before the fire. In particular, engineers could have reported the problems in obtaining a fuel tight seal during the periodic maintenance that took place immediately before the incident. Both of these apparent inadequacies can be described in terms of less than adequate inspection plans and less than adequate maintenance plans. Similarly, the failure to inspect the engine room between 15:15 and 16:00 can be characterised as the result of less than adequate procedures. It is important to reiterate that these are subjective interpretations of the failures that were identified during the first stage of the analysis. For instance, it could be argued that the failure to inspect the engine room between 15:15 and 16:00 was not simply the result of inadequate operating procedures. Better protection might have been offered if operators had been expected to document their inspection activities. This would have led the same Inspection LTA failure to have been classified under the Monitoring points LTA node of the Why? branch. Similarly, it can be argued that the inspection failure was due to inadequate training about the importance of these activities. This, in turn, could be due to a managerial failure to identify such a training requirement; Why?:Risk Assessment LTA: Safety Program Review LTA: Design and Development Plan LTA: Operational Specification LTA: Training LTA. Alternatively, it might be argued that the lack of inspection was not due to any of these factors but to management's failure to motivate staff to perform necessary safety inspections: Why?:Risk Assessment LTA: Safety Program Review LTA: Design and Development Plan LTA: Operational Specification LTA: Motivation LTA. These observations illustrate a number of important points about causal analysis using the MORT approach. Firstly, a number of different causal factors can be associated with the items identified in the first stage of the analysis. Some of these factors are not mutually exclusive. So, for example, inadequate inspection procedures might be compounded by a lack of monitoring points. Even if inspection procedures had been well-defined, motivational problems can 'dissuade' individuals from effectively following monitoring requirements.

Secondly, the Nanticoke case study supports a number of important observations about the nature of any causal analysis. It is difficult to be certain about which causal hypotheses, the nodes of the Why branch in the MORT diagram, can actually be applied to this incident. The available reports and documentation provide very little information about the motivation of the crewmembers or about the written procedures that were available to key personnel. Further investigations would, therefore, be necessary before any conclusions could be reached about these potential causes. An important strength of the MORT approach is that it directs investigators towards these potential hypotheses that must then be supported by further investigations. This offers a strong contrast to many event-based approaches. There is often an implicit assumption that counterfactual reasoning over a temporal model of event sequences can provide sufficient information about the underlying causes of an incident. This is a strong assumption. Chapter 10 has shown how NASA and the US Department of Energy have partially addressed these concerns by recommending the use of Tier or Compliance analysis to supplement the counterfactual reasoning afforded by ECF modelling.

The other control failures identified in Table 11.6 can be analysed in a similar fashion. Inadequate maintenance was recognised by the manner in which modifications to the forward filter cover and bolt sealing surface left grooves that made it hard to achieve a fuel-tight joint. This can potentially be explained in terms of inadequate maintenance and inspection plans under the path Why?:Risk Assessment LTA: Safety Program Review LTA: Design and Development Plan LTA. Operability problems including the lack of any new gaskets and the problems associated with the reuse of deformed gaskets can be associated with a management failure to conduct an adequate hazard analysis. Supervisory problems such as the failure to identify and correct damage incurred during previous modifications to the filter can be interpreted as the result of inadequate procedures. For example, a fault reporting system might have altered the chief engineer to the watchkeeping engineer's problems in achieving a sufficient seal on the filter. The failure to monitor the engines adequately between 15:15 to 16:00 can be interpreted as a failure of supervision in the operational specification of the system. The lack of any emergency exit forced the chief engineer to follow hand rails out of the engine room. This can be seen as a failure in the design basis of the ship as it was being operated immediately before the incident. Additional emergency lighting might, arguably, have supported the chief engineer's exit from a hazardous situation. Finally, the failure to prevent a recurrence of four previous engine fires on Canadian ships within six months of the Nanticoke incident can be associated with a failure to review the overall safety programme over previous years. In particular, Transportation Safety Board of Canada argued that previous warnings, such as that contained in Ship Safety Bulletin 13/85, had not resulted in adequate barrier being placed between potential fuel sources and adjacent exposed, hot surfaces.

Sub-Tree: Management Less Than Adequate (LTA)

Why?	What?	Description
Risk Assessment LTA	Hazard	Danger of a fire being started by a pressurised fuel release from a fuel filter onto the adjacent indicator tap or uncovered exhaust manifold.
	Target	People and systems in the engine room. Everyone on the vessel. Other ships in the same area as the Nanticoke.
Hazard Analysis LTA	Control: Operability problems	No new copper gaskets. Copper gaskets are deformed by use and pose more problems in obtaining a tight seal even if they have been annealed.
Inspection Plan LTA	Barrier: Did not use	More frequent inspections might raised the alarm sooner.
	Control: Inspection LTA	More regular inspections of filter assembly might reduced likelihood of fire. Crew members might have noticed the leak before ignition.
	Control: Inspection LTA	Engine room was not inspected between 15:15 and 16:00.

Table 11.7: MORT (Stage 2) Analysis Form

Tables 11.7 and 11.8 summarise the findings from the second stage of our MORT analysis. As can be seen, each of the nodes from the why branch in the MORT diagram can be represented as a row in the table. The what nodes that were identified during the first stage of the MORT analysis are then listed next to each of the why nodes if the corresponding (managerial) failures are perceived to have caused the more immediate failures that contributed to the incident. For example, the lack of adequate monitoring points to encourage compliance with inspection procedures is seem to have been a cause of the crews failure to adequately inspect the engine room between 15:15 and 16:00. It is important not to underestimate the significance of such tables. As mentioned, they provide a focus for continued discussion and analysis amongst the members of an investigation team.

The MORT analysis forms, illustrated in Tables 11.7 and 11.8, also act as a focus for other forms of analysis. For instance, the US Department of Energy have argued that investigators can sum the number of *what* factors associated with each why node to provide 'a measure of how widespread the element inadequacy is'. [204] In Tables 11.7 and 11.8 this can be done by counting the number of rows for each why? node. This would yield the following rankings for the Nanticoke case study. Inspection plan LTA is the only causal factor that is associated with three specific what failures. Risk Assessment LTA, Maintenance Plan LTA and Design Basis LTA are all associated with two specific failures. Hazard Analysis LTA, Monitoring Points LTA, Procedures LTA, Supervision LTA and Safety Program LTA are identified as the causes of a single failure in the accident/incident branches of the MORT diagram.

A number of objections can be raised to this form of analysis. The subjective nature of both stages in the MORT method can create considerable differences in the results that are obtained from this simple summation of accident factors. Similarly, it can be argued that different weights should be associated with each of the causal factors in the why branch of the MORT diagram. For instance, investigators may identify numerous instances in which operating procedures were inadequately specified. Changes in equipment design, in the operating environment and in regulatory requirements can prevent even the most assiduous operator from ensuring that all operating procedures are correctly documented. It might, therefore, be argued these problems are not as serious as less numerous maintenance failures. For instance, the Nanticoke incident might have had far worse consequences had the Halon system not been available to the Captain once his fire-fighting teams had been beaten back. Rather than develop more complex procedures for deriving aggregate weightings from MORT analysis form's, we adopt the more usual practice of assuming that investigators will use their skill and expertise to determine the overall significance of each row within Tables 11.7 and 11.8.

Previous paragraphs have described how the first stage of MORT analysis identifies what occurred during an incident. The second stage goes on to identify causal factors by asking why these failures arose. We have not, however, described the process by which root causes might be distinguished from the wider causal factors to the right of the MORT diagram. Several authors have argued that the concept of a 'root cause' originates with Johnson's early work on MORT [430, 444]. For example, Briscoe developed an analytical technique in which root causes are literally represented by the roots of the MORT diagram [96]. Investigators simply trace the more detailed why factors, identified in the Analysis Forms of Tables 11.7 and 11.8, up through the tree to identify the higher-level branches that represent the wider causes of managerial failure. The following list summarises the main categories that were identified by Briscoe's root cause analysis technique. Most of the items are relatively straightforward. Bridge elements represents the manner in which high-levels of management implement safety-related management policies throughout the various intermediate tiers of management within an organisation.

1. Policy
2. Policy Implementation
   - Line/staff responsibility
   - Accountability
   - Vigour and example
   - Methods and criteria analysis
3. Risk assessment
   - Safety-information systems
   - Hazard-analysis process
   - Safety-programme audit
4. Bridge elements
   - Management services
   - Directives
   - Budget
   - Information flow

Many of the causal factors that were identified for the Nanticoke case study can be broadly grouped under the 'hazard analysis process' root cause. Management failed to appreciate the dangers of the maintenance and inspection practices that were identified in Tables 11.7 and 11.8. Alternatively, if those dangers were recognised then it can be argued that there was an inadequate safety-programme audit because such practices were permitted to continue even after warnings such as that contained in Safety Bulletin 13/85.

Briscoe's approach is not the only checklist form of root cause analysis that might be applied after the second stage of a MORT analysis. For example, the International Loss Control Institute have developed a model of incident causation that extends the domino theory [85]. This approach proposes a number of further root causes in addition to those proposed by Briscoe [444]. These focus on common reasons behind failures at the workplace level:

Sub-Tree: Management Less Than Adequate (LTA)		
Why?	What	Description
Maintenance Plan LTA	Control: Inspection LTA	More regular inspections of filter assembly might reduced likelihood of fire. Crew members might have noticed the leak before ignition.
	Control: Maintenance LTA	Modifications to forward filter cover and bolt sealing surface left grooves that made it hard to achieve a fuel-tight joint.
Monitoring points LTA	Control: Inspection LTA	Engine room was not inspected between 15:15 and 16:00.
Design basis LTA	Barrier: Did not provide	Fire burnt through common cable tray containing all of the steering systems. Nanticoke was disabled until alternative power source was rigged for steering gear.
	Emergency actions LTA	No emergency exit from the control room. Chief engineer had to follow hand rails out of the engine room.
Procedures LTA	Control: 1st Line Supervision LTA	Failure to identify and correct damage incurred during previous modifications to the filter.
Supervision LTA	Control: 1st Line Supervision LTA	Failed to monitor engines during interval prior to the fire (15:15 to 16:00).
Safety Program Review LTA	Did not prevent 2nd accident	Four similar engine room fires occurred on Canadian ships within six months of the Nanticoke incident. Ship safety bulletin (13/85) had not resulted in adequate barrier being placed between potential fuel sources and adjacent exposed, hot surfaces.

Table 11.8: MORT (Stage 2) Analysis Form Continued

1. inadequate health and safety programme
2. inadequate health and safety programme standards
3. inadequate compliance with health and safety programme standards

Further additions might be made. Chapter 3 argued that the regulatory environment has a profound impact upon managerial behaviour. The decision only to apply the revised wiring requirement to vessels built after 1st September 1984 left the Nanticoke in a particularly situation when the fire burnt through the common cable tray that contained all of the steering systems. The decision to include such regulatory influences as a potential root cause within a MORT diagram depends upon the position of the investigator within an incident reporting system. In some schemes, typically those run by independent reporting agencies, it is possible for investigators to address these more general issues that might otherwise lie outside the scope of a conventional MORT analysis. If investigators decide to introduce regulatory and workplace factors, mentioned above, then these factors must appear as potential root causes in the upper levels of a revised MORT diagram. This increases the scope of the root cause analysis. Investigators must, however, navigate an increasingly complex diagram to identify those leaf nodes that best describe why particular failures occurred.

The MORT approach offers a number of significant benefits. In particular, it provides an early example of the way in which an engineering approach to safety, typified by barrier analysis, can be combined with broader managerial concerns. This blend of concerns has provides detailed insights into the way in which particular management activities contribute to many accidents and incidents [764]. The distribution and delegation of responsibility without adequate supervision often emerges as a common theme in MORT analyses. Similarly, the failure to implement well-specified safety plans can also be identified as a recurring pattern. There remains a considerable debate about whether or not these recurring themes are artifacts of the MORT analysis or whether they reflect common problems for different safety-critical systems [346]. A number of authors have, however, proposed automated tools that might automatically detect such recurring causal patterns amongst a 'database' of incident reports [457].

MORT offers a number of further benefits that relate more narrowly to the management of any investigation. The elements of the diagram direct investigators towards the potential causes of an incident. This helps to ensure that analysts consider a broad range of causal factors. The use of the tree can also provide necessary guidance for inexperienced investigators. It provides a common structure and format that encourages consistency in the investigatory process. The method associated with the tree is intended to ensure that investigators consider both what happened and why the incident occurred. The use of tabular check lists helps to communicate the products of a causal analysis to others within an investigatory team. Finally, the summary data that can be obtained from MORT tables, such as that illustrated in Table 11.5, can be used to monitor the changing causes of incidents across different geographical regions or organisational boundaries.

A number of limitations also restrict the utility of MORT as a tool for the causal analysis of safety-critical incidents. In contrast to STEP, this approach best be applied once investigators have already obtained a significant amount of information about an incident. Some proponents have argued that incident modelling, using ECF or accident Fault Trees, should be a prerequisite to any MORT analysis. In this view, counterfactual reasoning is used to identify causal factors that are then classified using the what branch of the tree. Instead of using Tier or Non-compliance analysis as in Chapter 10, investigators can then apply MORT to classify root causes against the why branch. Unfortunately, the perceived complexity of the MORT diagram and the potential overheads of such an integrated approach have dissuaded many analysts from exploiting these techniques [486].It is, therefore, seldom used in its full form without regulatory sanction. Munson argues that MORT is used more as a pro-active tool to support the analysis of a safety-critical design than it is as an accident investigation technique. This is due to the "nature of the nuclear industry, identifying possible loopholes in the safety system to eliminate hazards is more cost effective and publicly expedient than after the accident occurs" [552].

The leafs of MORT and mini-MORT diagrams may not capture the specific causes of an incident [290]. This should not be surprising. These diagram reflect the inevitable trade-off between large and unwieldy structures that embody many causal distinctions and more compact trees that provide a smaller number of more generic categories. As we have seen, investigators can extend MORT diagrams to address these limitations. This can, however, create inconsistencies within an incident reporting system. For instance, other investigators may not have used the new category in previous investigations. The extension of the MORT diagram can also create external inconsistencies between incident reporting systems if other organisations choose not to exploit the amended MORT diagram. Such problems can dissuade investigators from searching for causal factors that are not represented on the MORT diagram.

11.2.2 Prevention and Recovery Information System for Monitoring and Analysis (PRISMA)

As we have seen, ECF, MES and STEP help analysts to reconstruct the event sequences that contribute to incident and accidents. Different forms of counterfactual reasoning can then be used to distinguish between the causal factors and contextual details that are represented in these incident models. These techniques all focus on the specific events that occurred during a particular incident. Investigators must use a range of complementary approaches, such as Tier analysis, to identify the more generic root causes from the results of these more focussed techniques. In contrast, MORT relies upon investigators already having a relatively detailed understanding of the particular events leading to a mishap. The associated diagram and tabular form can be used to classify specific causal factors into a number of more general categories. It is important not to underestimate the significance of this distinction between MORT and the previous techniques. ECF, MES and STEP focus on 'singular causality' [677]. MORT focuses on the notion of 'general causality' that was introduced in Chapter 7.

A number of researchers have recognised the distinctions between particular and general causality that are embodied within ECF, MES, STEP and MORT. They have responded by developing more integrated approaches that are intended to support both the reconstruction of the specific events that lead to an incident and the identification of more general causal factors. The Prevention and Recovery Information System for Monitoring and Analysis (PRISMA) is one example of this dual technique [840, 841]. This approach is also different from those introduced in previous sections because it was specifically developed to enable organisations to monitor and respond to incident reports. It was not intended to support accident investigation.

Van Der Schaaf's motivation in developing PRISMA was to support the development of a quantitative database of incident data. This resource was to guide the detection and prevention of structural problems rather than the particular characteristics of individual incidents [844]. The PRISMA approach consists of three principle stages. The following paragraphs describe each of these stages and illustrates how they can be used during a causal analysis of the Nanticoke case study:

1. *Reconstruct the incident using a causal tree.*

2. *Use a classification model to identify generic factors.*

3. *Apply a classification/action matrix to identify potential counter-measures.*

Causal trees are similar to the Fault Trees that were introduced in Chapter 10. The overall structure of the tree reflects the chronology of an incident. The left-most branches indicate latent conditions or failures that occur relatively early in the course of events. The right-most branches are, typically, used to model recovery actions and interventions that mitigate the consequences of an incident. It is important to note, however, that causal trees are constructed using AND gates. Investigators must avoid the uncertainty that is implied by disjunction. Van Vuuren notes that "the main difference between a causal tree and a fault tree is that the top event in a causal tree is not a class of events but one particular incident, which actually occurred and for which the chain of causation can be discovered" [844]. In contrast to the MORT diagram, causal trees are intended to capture the 'who', 'what' and 'where', they do not explain 'why' an incident may have occurred. Figure 11.8 presents a causal tree for the Nanticoke case study.

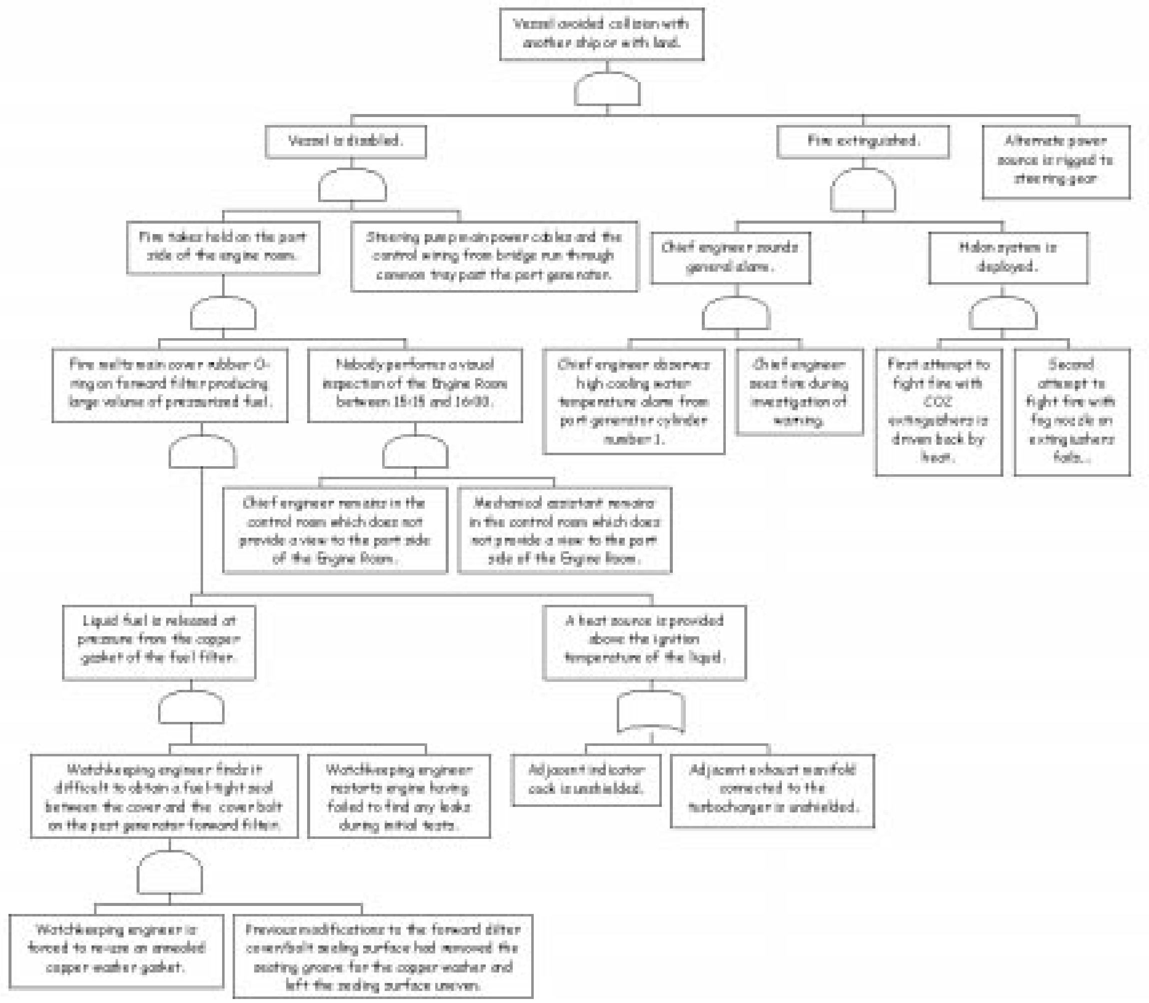
Vessel avoided collision with another ship or with land.
Vessel is disabled.
Fire extinguished.
Alternate power source is rigged to steering gear
Fire takes hold on the port side of the engine room.
Steering pump main power cables and the control wiring from bridge run through common tray past the port generator.
Chief engineer sounds general alarm.
Halon system is deployed.
Fire melts main cover rubber O-ring on forward filter producing large volume of pressurized fuel.
Nobody performs a visual inspection of the Engine Room between 15:15 and 16:00.
Chief engineer observes high cooling water temperature alarm from port generator cylinder number 1.
Chief engineer sees fire during investigation of warning.
First attempt to fight fire with CO2 extinguishers is driven back by heat.
Second attempt to fight fire with fog nozzle on extinguishers fails..
Chief engineer remains in the control room which does not provide a view to the port side of the Engine Room.
Mechanical assistant remains in the control room which does not provide a view to the port side of the Engine Room.
Liquid fuel is released at pressure from the copper gasket of the fuel filter.
A heat source is provided above the ignition temperature of the liquid.
Watchkeeping engineer finds it difficult to obtain a fuel-tight seal between the cover and the cover bolt on the port generator forward filter.
Watchkeeping engineer restarts engine having failed to find any leaks during initial tests.
Adjacent indicator cock is unshielded.
Adjacent exhaust manifold connected to the turbocharger is unshielded.
Watchkeeping engineer is forced to re-use an annealed copper washer gasket.
Previous modifications to the forward filter cover/bolt sealing surface had removed the seating groove for the copper washer and left the sealing surface uneven.

Figure 11.8: A Causal Tree of the Nanticoke Case Study

There are considerable differences between causal trees and the various modelling techniques in ECF, STEP and MES. For instance, analysts can annotate the nodes in a causal tree with natural language labels that do not distinguish between events and conditions. These annotations are intended to capture observations about the course of an incident in a flexible and informal manner. One consequence of this is that it can be difficult for investigators to distinguish the actions of particular individuals during an incident. Rather than grouping these along a single row, as in STEP, they can be distributed across the many different nodes of a causal tree. The lack of typing information also means that ambiguities and omissions can weaken the integrity of these diagrams. For instance, some of the proponents of this approach have published trees whose nodes are labelled He was standing next to the person or he saw the falling object. Such annotations work well for small examples but cannot easily be maintained for more complex incidents, such as the Nanticoke case study. Figure 11.8, therefore, explicitly identifies the key individuals who were identified during the primary investigation into this mishap.

There are further differences between the causal tree of Figure 11.8 and the checklist approach embodied in MORT. In particular, the nodes represent both positive or mitigating factors as well as the failures that contribute towards an incident. As can be seen, this diagram denotes the way in which the chief engineer eventually noticed the high cooling water alarm from the port generator, cylinder number 1. It also records the successful use of the Halon system to extinguish the fire after two attempts to use carbon-dioxide extinguishers were beaten back by the heat. These right-hand branches are a significant strength of the PRISMA approach to incident modelling. As we have reiterated, organisation learning depends not simply upon recognising the causes of failure but also on promoting those actions that help to combat previous failures.

Every leaf nodes represents a causal factor. At first sight, this might appear to lack the sophistication of the more elaborate counterfactual approaches from previous sections. It is important to remember, however, that causal trees are entirely constructed from AND gates. It, therefore, follows that if any of the leaf nodes are not true then the top level incident will not be true. In consequence, this approach mirrors the counterfactual decision procedure of ECF, MES and STEP. There are, however, some exceptions to these general comments. As can be seen from Figure 11.8 it may still be necessary to include an OR gate within the causal trees that represent particular incidents. As with the Allentown explosion in Chapter 9, it is likely that we shall never be able to determine the exact ignition source for the Nanticoke case study. Transportation Safety Board of Canada investigators identified the indicator tap and exhaust manifold as potential sources. They were, however, unwilling to commit themselves to which was the most likely cause of the ignition. This uncertainty is denoted by the OR gate in Figure 11.8. As we shall see, this introduces a number of theoretical problems for the application of the PRISMA technique.

The second stage in the application of the PRISMA approach uses a classification model to associate a more generic root cause with each of the causal factors that are denoted by leaf nodes. This focus on the leaf nodes is justified by the observation that internal nodes are often the result or consequence of these other events and conditions. For instance, in Figure 11.8 two leaf nodes represent the facts that Previous modifications to the forward filter cover/bolt sealing surface had removed the seating groove for the copper washer and left the sealing surface uneven and Watchkeeping engineer restarts engine having failed to find any leaks during the initial tests. These two factors helped to create a situation that is represented by the interior node Watchkeeping engineer finds it difficult to obtain a fuel-tight seal between the cover and the cover bolt on the port generator forward filter. The re-use of the annealed copper gasket and the damage caused by previous modifications are seen to be causes of the engineer's subsequent difficulties. They are the focus for the subsequent classification rather than the interior node that represents the consequence of those two factors.

The second stage of the PRISMA analysis also, typically, focuses on the left-hand side of the causal tree. Recovery or mitigating factors are typically located on the right-hand side of the tree because they, typically, occur after the initiating conditions. These factors are important because they provide insights into protection mechanisms that successfully mitigated the potential consequences of an incident. For instance, the right-hand nodes of Figure 11.8 represent the crews actions that ultimately extinguished the fire on the Nanticoke. They also describe how an alternative power supply was rigged to the steering gear so that the crew could regain control of their vessel. These mitigating factors are not considered during this second stage of analysis. They represent remedial actions rather than causal factors. It is important to provide a procedure that can be used to distinguish causal factors from other mitigating actions in a causal tree. This can be done using the counterfactual reasoning that was introduced in previous paragraphs. For each node in a causal tree then investigators must ask whether the incident would have occurred if that node had not occurred. If the answer is no then the node represents a true causal factor and it is used during the subsequent classification. If the answer is yes then the node is not carried forward into any subsequent analysis. For example, the omission of a mitigating factor is likely to have exacerbated an incident rather than prevented its occurrence.

Unfortunately, the presence of disjunctions in a causal tree can considerably complicate this use of counterfactual reasoning. For example, if ask 'would the Nanticoke incident have been avoided if the adjacent indicator tap been shielded' then the answer would be no. The ignition might have been caused by the exhaust manifold. Conversely, if we ask 'would the incident have been avoided if the adjacent exhaust manifold had been shielded' then the answer would also be no. The ignition might have been caused by the indicator tap! Such problems can be resolved by further empirical studies or mathematical modelling. As we have seen in the Climate Orbiter case study, it is important not to over-estimate our ability to reconstruct the events leading to many incidents. The Nanticoke mishap is not the only case study in which such problems arise. For example, there are a number of competing hypotheses about the event sequences that led to the loss of the Deep Space 2 probes. In Chapter 10 we focussed on the potential problems that may have arisen during impact with the Mars surface. However, the probes may also have been damaged during separation from the cruise stage of the Polar Lander. If we ask 'would the incident have been avoided if the probes successfully separated from the cruise stage' then the answer is no. The probes might have been destroyed on impact with the planet surface. Conversely, if we ask 'would the incident have been avoided if the probes were resilient enough to survive impact with the planet surface' then the answer would again be no. Even if they had been capable of surviving the impact, they may not have reached that stage of the mission if problems had occurred during separation. Previous chapters have argued that such problems can be avoided by applying counterfactual reasoning over several different competing failure scenarios. In this view, investigators assume that one of the competing sets of events occurred. For instance, that the Nanticoke ignition was started by the adjacent indicator tap and not be the exhaust manifold. Counterfactual reasoning can then be applied as before. The lack of shielding can, therefore, clearly be identified as a causal factor. This reasoning process can then be repeated for the alternative failure scenarios. We term this counterfactual reasoning by *proxy*. Any ambiguity, such as that represented by the OR gate in Figure 11.8 is replaced by an assumed version of events. This assumption can then, in turn, be substituted by alternative event sequences during subsequent analysis. For instance, the assumption that the exhaust manifold provided the ignition source can be replaced by an assumption that the indicator tap helped to cause the fire.

The leaf nodes that represent causal factors in the Nanticoke case study are summarised as follows:

- Steering pump main power cables and the control wiring from the bridge run through a common tray past the port generator.

- Chief engineer remains in the control room which does not provide a view to the port side of the engine room.

- Mechanical assistant remains in the control room which does not provide a view to the port side of the engine room.

- Watchkeeping engineer is forced to re-use annealed copper washer gasket.

- Previous modifications to the forward filter cover/bolt sealing surface had removed the seating groove for the copper washer and left the sealing surface uneven.

- Watchkeeping engineer restarts engine having failed to find any leaks during the initial tests.

- Adjacent indicator tap is unshielded.

- Adjacent exhaust manifold is unshielded.

As mentioned, these causal factors are then categorised using a classification model that guides the investigators analysis. These models are used to associate more general root causes with the specific causal factors that are obtained from the causal tree. They can therefore be thought of as a further variant of the checklist approach, introduced in Chapter 10. PRISMA was initially developed to exploit the Eindhoven Classification Model, illustrated in Figure 11.9. This model was derived from an investigation of the causes of safety-related failures in the chemical process industry [840]. Since that time, however, a number of more detailed models have been developed to support the analysis of incidents in the medical and steel production domains [844]. For example, Figure 11.10 illustrates a medical classification scheme. The Eindhoven Classification Model focuses on three main categories of failure: technical; organisational and human. These can then be sub-divided into a number of more detailed causal factors. For instance, causal factors that relate to human behaviour can be associated with rule, knowledge or skill-based performance. These distinctions reflect Rasmussen's model of cognition introduced in Chapter 3. Similarly, organisational root causes are divided into inadequate operating procedures or ill-advised management priorities.

The classification process follows a fixed order [844]. Investigators must first determine whether the causal factor relates to the technical work environment. If the answer is yes, then the investigator must use the model in Figure 11.9 to determine the nature of that technical failure. Was the root cause related to an engineering, construction or materials problem? If the causal factor cannot be associated with a technical root cause then investigators must consider the organisational context of the incident. If technical and organisational factors are ruled-out then human behaviour can be considered as a root cause. This ordering is entirely deliberate. As with MORT, the detailed architecture of the classification scheme reflects the perspective and priorities of its developers. In this case, the Eindhoven Classification Model places human behaviour last so that investigators are forced to consider other causal factors before 'blaming' individual operator error.

The Eindhoven Classification Model from Figure 11.9 can be used to identify root causes from the causal leaf nodes of Figure 11.8. Table 11.9 summarises the results of this analysis. As we have seen, the use of a common tray to route all of the steering power and control wiring was identified as a causal factor in the loss of control that followed the fire. The decision to employ this approach can be associated with a technical failure in the engineering of the vessel. In consequence, Table 11.9 associates the wiring layout with the TE root cause from the Eindhoven Classification Model. The same categorisation can be applied to the manner in which previous modifications had removed the seating groove for the copper washer and left the sealing surface uneven. Previous sections have argued that this damage reflects incorrect maintenance procedures. It can, however, be argued that the removal of the seating groove was a consequence of previous maintenance problems. This again illustrates how the application of causal analysis techniques, such as PRISMA, are not an end in themselves. They raise questions that can only be resolved through further investigation.

Table 11.9 associates the same root cause with both of the hypothesised ignition sources. The lack of shielding around the indicator tap and exhaust manifold is associated with a technical failure in the construction of the engine assembly. It could be argued that these problems relate more to the engineering or design of the engine and filter rather than to its construction. This example also illustrates how distinctions that are meaningful within one industry need not be important in other domains. The differences between engineer, construction and materials are clearly defined within Van Der Schaaf's initial studies of the chemical process industries [840]. They are, however, less clear cut for our maritime case study. Such observations illustrate the need to derive classification models that capture pertinent root causes within a particular application domain.

It is also possible to challenge our claim in Table 11.9 that the re-use of the annealed copper washer gasket stemmed from a failure in organisational operating procedures. The re-use of copper gaskets that had previously been deformed under high operating pressures should not have been permitted. Conversely, it can also be argued that this failure stems more from a technical failure to ensure that the engineers were supplied with adequate materials. This illustrates the importance of both documenting the outcome of any root cause analysis and the associated justifications that support a particular categorisation. These documents can be shown to other investigators and safety managers to validate the products of any causal analysis. Any conflicts might be resolved

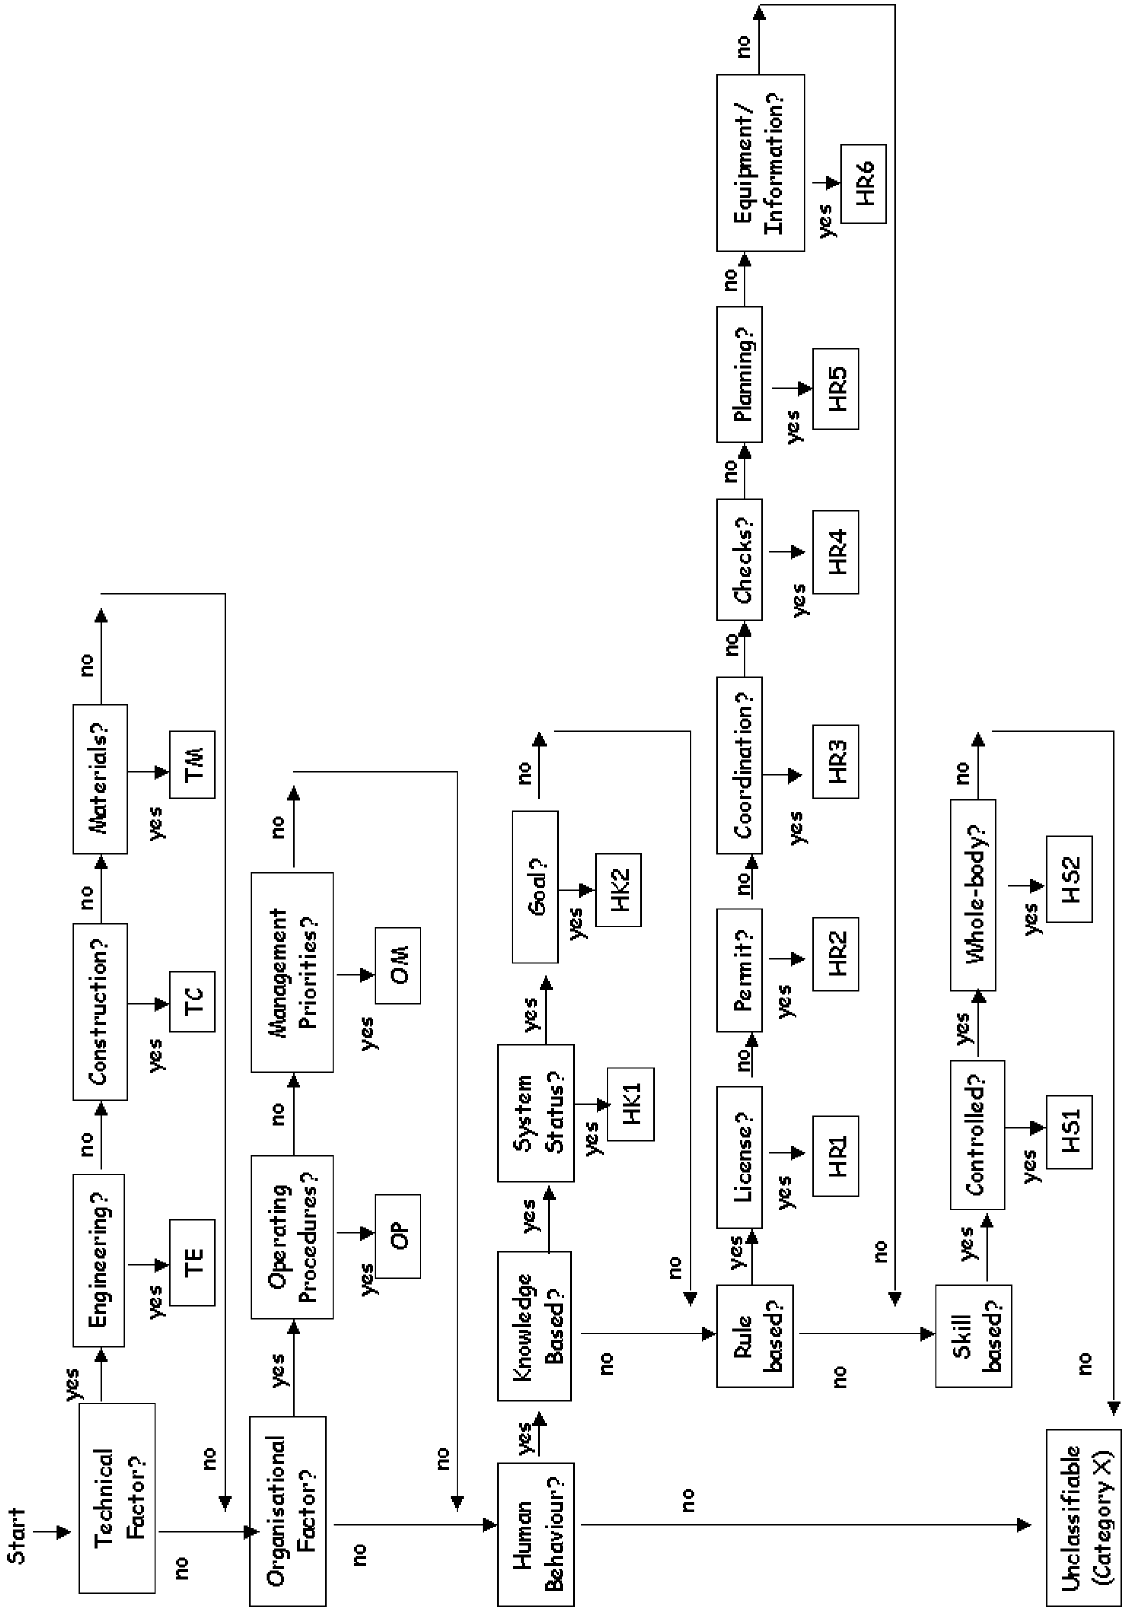

Figure 11.9: The Eindhoven Classification Model [840]

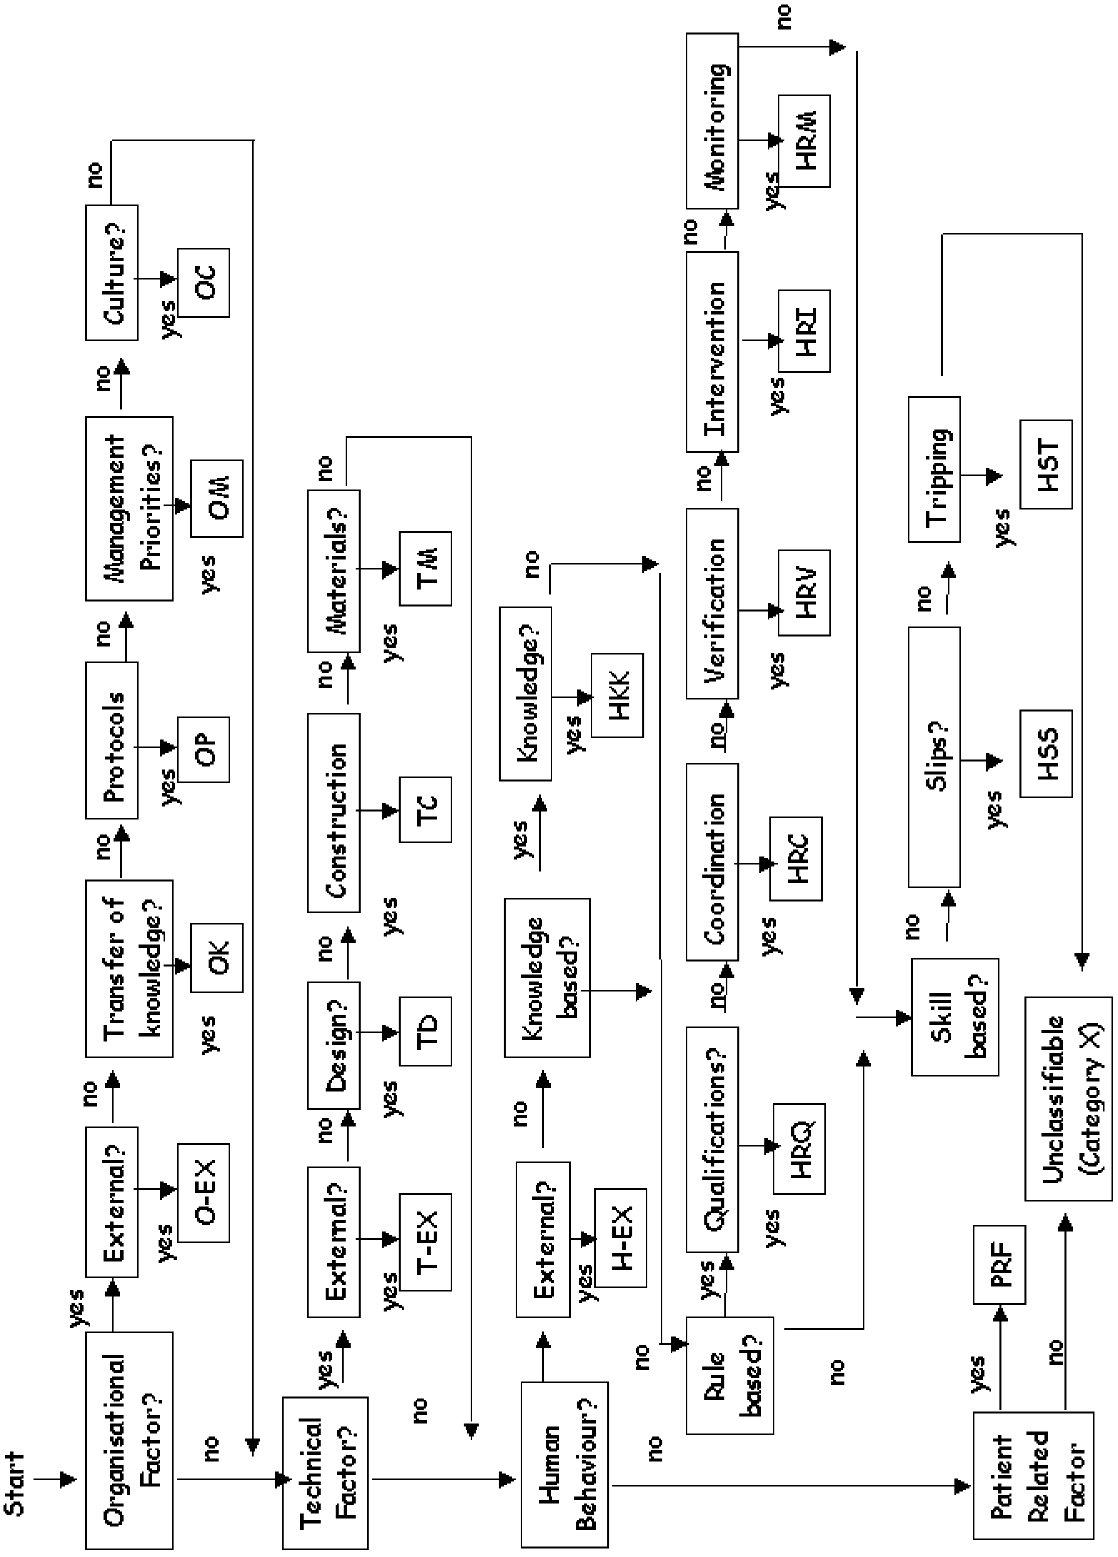

Figure 11.10: Classification Model for the Medical Domain [844]

Causal factor	ECM Classification
Steering pump main power cables and the control wiring from the bridge run through a common tray past the port generator.	TE - Technical Factor: Engineering.
Chief engineer remains in the control room which does not provide a view to the port side of the engine room.	HR4 - Human Behaviour : Rule Based : Checks.
Mechanical assistant remains in the control room which does not provide a view to the port side of the engine room.	HR4 - Human Behaviour : Rule Based : Checks.
Watchkeeping engineer is forced to re-use annealed copper washer gasket	OP - Organisational Factor : Operating Procedures.
Previous modifications to the forward filter cover/bolt sealing surface had removed the seating groove for the copper washer and left the sealing surface uneven.	TE - Technical Factor: Engineering.
Watchkeeping engineer restarts engine having failed to find any leaks during the initial tests.	HK1 - Human Behaviour : Knowledge Based : System Status.
Adjacent indicator tap is unshielded. Adjacent exhaust manifold is unshielded.	TC - Technical Factor: Construction.

Table 11.9: PRISMA (Stage 2) Summary Table

by encouraging analysts to associate multiple root causes with each of the causal factors that are identified during previous stages of analysis. This approach is not generally encouraged [844]. There is a danger that the unnecessary proliferation of root causes will hide important information about the factors that contributed to an incident.

Table 11.9 identifies a number of root causes that stem from human factors problems. The Chief engineer and the mechanical assistant remained in the control room from 15:15 to 16:00. They could not observe the port side of the engine room from this position and so failed to observe the fire as it began to take hold. This can be interpreted as a rule-based failure to perform necessary checks. It can be argued that the watchkeeping engineer's decision to restart the engine after failing to find any leaks was the result of a knowledge-based failure in their interpretation of the state of the system. Such findings must, however, be treated with caution. The Transportation Safety Board of Canada investigators provided very little information about the decision to restart the engine. It is, therefore, difficult to be certain of the root causes that may have influenced the engineer's behaviour. It was not anticipated that so many root causes would relate to human factors problems in our analysis of the Nanticoke case study. The ordering of the Eindhoven Classification Model considers there factors after other technical and organisational factors. The analysis may reflect the reliance upon human intervention in the Nanticoke case study. Our findings might also be unnecessarily biased by the evidence that was available in the aftermath of this incident.

The final stage in any PRISMA analysis is to identify 'recommended' actions that might address each root cause. PRISMA provides a classification/action matrix to support this task. These tables link each category of the classification model to a ranked list of interventions. These responses are ordered according to their perceived cost effectiveness. They may relate to improved acquisition or equipment design, to better procedures, information management or communication, to revised training practices or motivational activities [840]. The exact nature of the table will vary from industry to industry and from organisation to organisation. The effectiveness of particular recommendations can be affected by the wider safety culture in a company. It can also be influenced by the financial and other resources that are available to an investigator. In consequence, the entries in a classification/action matrix are likely to change over time. Safety reviews are liable to identify new rankings for the effectiveness of particular recommendations.

The classification/action matrices represent an important aspect of PRISMA that has not been addressed by the other causal analysis techniques in this Chapter. ECF, MES and STEP focus on the events leading to an incident or accident. MORT does provide means of analysing the response to an incident. Recommendations from any previous incidents should ensure that an oversight or omission becomes an assumed risk. These "are defined as only those risks that have been analysed and accepted by the proper level of management; unanalysed or unknown risks are not considered to be Assumed Risks" [203]. None of these techniques provides explicit means of ensuring a consistent response to similar incidents. Not does it provide means of monitoring the effectiveness of that response.

	Organisational Factors				
	External Factors (O-EX)	Knowledge Transfer (OK)	Operating procedures (OP)	Manag. priorities (OM)	Culture (OC)
Inter-departmental communication	X				
Training and coaching		X			
Procedures and protocols			X		
Bottom-up communication				X	
Maximise reflexivity					X

Table 11.10: Example PRISMA Classification/Action Matrix [844]

Table 11.10 illustrates the general format of the classification/action matrices that are advocated by the PRISMA approach. This particular example is derived from the medical classification model. The increased number of organisation categories in this model provides an interesting insight into the nature of medical incidents when compared with the abridged version in the original Eindhoven model, illustrated in Figure 11.9 [844]. As can be seen in Table 11.10, incidents that involve a failure in knowledge transfer within an organisation might result in revised training and coaching practices. Failures that stem from problems involving operating procedures will, as expected, result in revised procedures and protocols. The precise nature of such tables is determined by the context in which any recommendations will be applied. Individual organisations may also be forced to increase the level of detail that is represented within Classification/Action matrices such as that shown in Table 11.10. For example, a recommendation to improve training and coaching is not at a sufficient level of detail to encourage confidence that any recurrence will be avoided. The motivation behind this technique is summarised by Van Vuuren who argues that:

> "However, the incident data clearly shows decreased risk awareness and vigilance as main contributors to adverse group behaviours, leading to incidents. Therefore, an organisation should reflect on its safety experiences and try to learn as much as possible from them. The correct level of risk awareness and vigilance can be maintained by reporting and analysing the often abundantly available near misses. Based on these analyses, feedback to the organisation can be provided to show the dangers of day to day practice. This way, a continuous circle of learning from its own safety experiences and measuring the safety performance of the organisation results." [844]

It is, however, possible to apply elements of Table 11.10 to the Nanticoke case study. Previous stages of the analysis argued that the re-use of the annealed copper washer gasket stemmed from a failure in organisational operating procedures. The re-use of copper gaskets that had previously been deformed under high operating pressures should not have been permitted. As might be expected, Table 11.10 suggests that this root cause might be combatted by revising the procedures and protocols that govern current maintenance practices.

A number of limitations can be identified for the PRISMA technique. Some of these relate to particular features of this approach, others are more general criticisms of checklist techniques. PRISMA, like MORT, offers greatest support after primary and secondary investigations have been completed. It depends upon investigators being able to construct the causal trees that have been illustrated earlier in this section. The proponents of PRISMA do, however, urge that the application of this approach should be based around critical incident interviews based on a technique developed by Flanagan in the 1950's [250]. This interview technique encourages individuals to describe situations in which the success or failure of an operation was determined by specific causes. It is argued, by extension, that the same approaches can be used to elicit information about mishaps for which the causes are less certain. This utility of this elicitation technique has been validated by considerable fieldwork. It also integrates well with the generation of causal trees that are intended to capture both good and poor performance. Critical incident interviews can, however, only provide part of the evidence that is necessary for the causal analysis of complex, technological failures. For example, it is unclear how information from automated logging systems or from regulatory documents might be integrated into these 'anecdotal' accounts. Similarly, there is little guidance about how to address the increasing complexity of many near-miss incidents, which involve individuals and systems from many different organisations and working groups.

The practical application of PRISMA has been assessed in a number of studies. For example, this approach has been used to identify the root causes of incidents from NASA's Aviation Safety Reporting System [529]. Investigators were trained to use a variant of the Eindhoven Classification Model. They were then asked to independently classify the same group of incident summaries. The intention was to assess interrater reliability using the PRISMA method. The results indicated that subjectivity might be less of an issue than has been claimed for checklist approaches, however, the investigation raised more questions than it addressed. More interestingly, this study identified a number of fundamental misconceptions that arose when investigators were trained to apply the PRISMA technique. For example, one participant in the trial was unhappy that they were able to provide an unambiguous classification for all of the incidents that were studied. They then went back to the dataset until they could classify some incidents under the X - unclassifiable category. Such incidents are instructive for a number of reasons. Firstly, they point to the difficulty of training investigators to use even simplified forms of the existing analytical techniques. Secondly, they point to the way in which individual differences can influence the successful application of these techniques. None of the other participants expressed this concern that some incidents should not be classified by the existing model! It is important to emphasise that these concerns are not simply centred on the PRISMA approach but can potentially affect all of the analytical techniques described in this book.

It is also possible to identify a certain confusion about the distinction between causal factors and root causes in the PRISMA technique. Van Vuuren has argued that root causes can be identified as the leaf nodes in the left-hand branches (i.e., the non-mitigating branches) of a causal tree [844]. In his view, classification model simply provide a means of grouping these root causes into categories that are amenable to statistical analysis. Managers can use the results of the classification process to monitor, for instance, how many incidents are caused by problems with operating procedures in a given time period. This is an interesting approach because, in some ways, it is the antithesis of MORT. Root causes are represented by the upper nodes of the MORT diagram. In Van Vuuren's view of PRISMA, root causes are denoted by the lower leaf nodes of a causal tree. The difference becomes apparent if we compare the leaf node Steering pump main power cables and the control wiring from bridge run through common tray past the port generator from the PRISMA causal tree with the corresponding *why* branch from the MORT analysis Barrier: Did not provide. As can be seen, the MORT approach more closely resembles our requirement that root causes should be more general than the causal factors that characterise a particular incident. In consequence, the previous pages have adopted the convention of referring to the non-mitigating leaf-nodes of a causal tree as *causal* *factors* and the elements of a classification model as *root causes*.

A number of general criticisms can be made about checklist approaches such as PRISMA and MORT [444].Previous paragraphs have already argued that investigators may be dissuaded from searching for potential root causes that do not appear on a checklist. Further biases can affect the selection of items within a checklist. For instance, items at the top or the bottom of a list are more likely to be selected than those in the middle [457]. Similarly, if certain classes of causal factors occur more frequently in a checklist then there is an increased likelihood that those factors will be identified. MORT provides an extreme example of this in which all root causes can be linked to management failures, neglecting regulatory, environmental or other workplace factors. Kjellén has also argued that investigators and supervisors are more likely to choose those causal factors on a checklist that are difficult to verify or that involve limited management responsibility [444]. There is an increased tendency to select factors that relate to individual failures or to adverse factors that are 'beyond the control' of senior and middle management. This partly explains MORT's bias towards managerial factors.

Checklist approaches also suffer from the wide range of biases that have been noted in previous chapters. Attribution errors make it more likely that investigators will select transient or environmental causal factors if they are implicated in an incident [444]. This is less likely to occur when investigators belong to an independent investigation agency. We have also seen how the lack of event-based models can also create problems for checklist-based approaches. techniques such as ECF analysis, MES and STEP provide a map of events that can be used to trace the development of an incident over time. If additional evidence becomes available then this can be directly used to revise these temporal models. In contrast, it can be more difficult to trace the impact of new information on the causal analysis supported by checklist techniques. Information about particular events can be distributed throughout the stage 1 and stage 2 tabular forms that support any MORT analysis. Similarly, it can be difficult to reconcile the temporal and causal relationships that are embedded within the gates of a causal tree.

11.2.3 Tripod

Previous sections have reviews a number of techniques to support causal reasoning about adverse occurrences. None of these techniques has, however, explicitly recognised the distinctions between catalytic and latent failures that has been emphasised in previous chapters. In contrast, the Tripod techniques were deliberately developed to account for this important distinction. The Tripod research project started in 1988 from a collaboration between the Universities of Leiden and Manchester. This collaboration has produced a range of analytical techniques. Tripod-Delta supports the predictive analysis of potential causal factors without the need for accident and incident statistics. Tripod-Beta provides more focussed support for incident and accident investigation [701]. The underlying analytical techniques have been widely used within the petrochemical industry [372, 853]. It is important to emphasise, however, that Tripod is not simply an accident or incident analysis technique. It's proponents argue that it offers a coherent philosophy based on the precept that safety management is essentially an organisational control problem.

Figure 11.11 sketches the model of incident and accident causation that underpins the Tripod method. It also illustrates the three key concepts that motivate the use of the name Tripod. Incidents and accidents provide important information about underlying, or root causes, of systems failure. These underlying or latent conditions are referred to as General Failure Types. As we shall see, they stem from the organisational, managerial and regulatory practices that create the preconditions for failure. The final leg of the tripod is provided by the active failures or unsafe acts that trigger an incident. These unsafe acts initiate hazards that can be mitigated by the proper use of barriers or may ultimately develop to compromise the safety of the target [205].

Tripod also provides a framework for thinking and for measuring the disturbances that affect safe operations. This measurement is based upon the General Failure Types mentioned above. These have strong similarities to the branches in a classification hierarchy, such as a MORT diagram or the Eindhoven Classification Model. There are also important differences. For instance, Tripod-Delta's measurement of potential disturbances to safe practice does not rely upon incident or accident

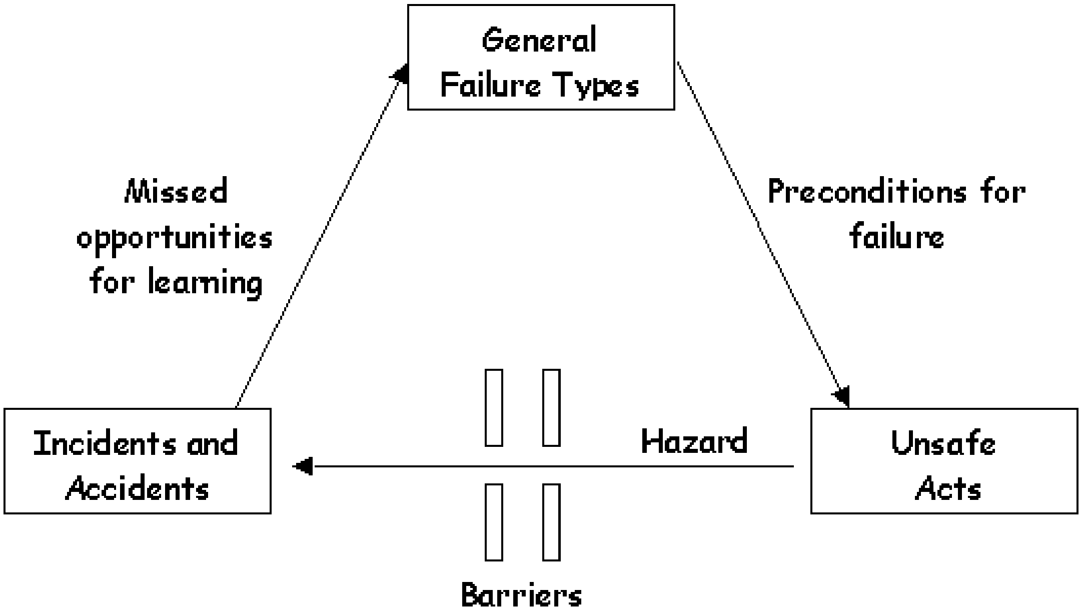

Figure 11.11: The Three Legs of Tripod

statistics. This contrasts strongly with the US Department of Energy's proposal to derive aggregate values for the root causes identified by MORT analysis.

Tripod relies upon an underlying model of causation. This assumes that incidents are caused by local triggering factors that combine with more latent General Failure Types. It is also assumed that organisations can do little to predict or address these local triggering factors. Reason uses the analogy that they are like mosquitos [701]. It does little good to swat at them individually; it is far better to drain the swamp in which they breed. In this case, the swamp represents the latent General Failure Types. These can be summarised by the following list. Each item emerged through close study of previous incidents rather than through any explicit empirical investigation. It should also be stressed that some failure types have consequences that promote other 'knock-on' failure types. For instance, inadequate maintenance management can lead to working conditions that increase the likelihood of operator error:

1. *Hardware.* Unsafe acts often result from the provision of inadequate equipment and materials. This can be the result of poor stock control, of problems in the supply chain, of component defects etc.

2. *Maintenance management.* Unsafe acts may also stem from the management rather than the execution of maintenance activities. For example, an incident may occur because necessary maintenance work was delayed or postponed.

3. *Design.* Unsafe acts can occur if designers fail to provide operators with sufficient information about the purpose and reliability of a device. Similarly, designers may provide inadequate information about the range of safe interventions that can be made with a device. They may also provide users with insufficient feedback about the state of a device.

4. *Operating procedures.* Unsafe acts may stem from procedures that either could not be applied in a given context or which contained dangerous advice or which contained advice that could not physically be complied with by an operator. Procedures may also be ambiguous in their application and in the guidance that they offer.

5. *Violation-inducing conditions.* Unsafe acts can stem from workplace or environmental pressures that encourage violations or discourage compliance. These factors may also promote erroneous behaviour, for instance, by exposing operators to hostile working conditions.

6. *Housekeeping.* Many incidents are caused by failures that are well known but which have not been adequately addressed over a long period of time. For example, problems in maintenance management can lead to hardware problems that become compounded over time.

7. *Incompatible goals.* Incidents can occur because individuals may be preoccupied or have goals that conflict with those that are intended to ensure the safety of the system in which they operate. The goals and working practices of groups can conflict with those of others within an organisation. Finally, there may be conflict between organisational objectives, such as profit or public approval, and safety.

8. *Communication.* Mishaps can be the result of system failures that impair communications channels. They can also stem from lost signals even when it is physically possible to transmit a message. For example, a safety warning might be delivered to the wrong person within an organisation. Even if a messages is successfully received, it can be misinterpreted or may arrive too late to ensure the safety of an application.

9. *Organisation.* Organisational structures can prevent individuals from responding to the lessons provided by safety-related incidents. For example, there may be divided responsibilities or conflicts of interest.

10. *Training.* Mishaps can occur if personnel lack the competence required to complete necessary tasks. This can occur if training is inadequately prepared, if it is curtailed, if it is not validated as providing the necessary instruction etc.

11. *Defence planning.* Mishaps can also occur if there are deficiencies in the detection, mitigation and remedial actions that are taken in the aftermath of an incident.

In common both with MORT and several other checklist approaches [385], General Failure Types stem from management decisions. Within each of these General Failure Types it is possible to distinguish two different levels of cause. Functional failures stem from decisions made by line managers, by designers, by planners etc. In contrast, source failures refer to more strategic decisions at senior management level. This has some similarities to the broad categories within Tier Analysis, described in Chapter 10.

As mentioned above, Tripod-Delta can be used in a pre hoc manner. It does not depend upon incident or accident statistics. This is important because, as we have seen, the insights that are provided by these information sources can be marred by under-reporting or by analytical biases. Reason argues that domain and task specialists can devise questions that will test for the presence of different General Failure Types before an incident occurs. For example, workers on an offshore platform might be asked 'was this platform originally designed to be unmanned?' or 'are shutoff valves fitted to a height of more than two meters?' [701]. These questions are intended to elicit highly focussed responses that are indicative of the more general General Failure Types, listed above. Software support has been developed to help administer these questionnaires. Approximately, twenty indicators are identified for each of the eleven General Failure Types. Once operators have completed these questions, the system compiles a bar chart that represents a Failure State Profile. This bar chart lists the General Failure Types according to the number of 'incorrect' questions that were answered by the operator. For example, the system asks twenty questions that relate to each General Failure Type. If eleven of the questions about hardware failures raised a potential cause for concern but only six of the questions about communication were answered in this way then hardware might be interpreted as a greater priority than communications issues. This represents a relatively crude interpretation of the analysis. It is recommended that the software be used three or four times a year and that any consequent decisions are based on trends rather than one-off values. For example, if we assume that an operator answered ten if ten out of twenty answers that the operator provided about hardware failure indicated that this was a significant cause for concern then this

General Failure Type would be ranked above any other failure types that The key point in all of this is that the questions, or indicators, help to trace the symptoms of a problem. The General Failure Types capture the underlying causes of future safety problems.

Tripod-Delta provides a general tool that can be used without incident and accident statistics. In contrast, Tripod-Beta was developed to provide incident analysis tools that can be used as an investigation progresses [216]. This explicit intention to support the investigatory process is similar to the motivation behind event cards in STEP. It contrasts sharply with the assumption in techniques, such as MORT or PRISMA, that the investigatory process has been largely completed. The Tripod-Beta software provides investigators with guidance about the elicitation process. As might be expected, investigators are prompted to go beyond the local triggers to identify latent General Failure Types. Hence, Tripod-Beta was deliberately intended to be compatible with Tripod-Delta.

Tripod-Beta analysis exploits many of the concepts that were introduced during the discussion of Barrier Analysis in Chapter 10. Investigators begin by identifying the targets that were affected by a potential hazard. They then have to trace the manner in which individual barriers were compromised during an incident. This is, typically, done by constructing a form of causal tree. At the root of the Tripod-Beta tree is an active failure that helped to compromise one of the barriers, mentioned above. The second level of the tree describes preconditions that had to be satisfied in order for the active failure to occur. For example, Figure 11.12 uses Tripod to analyse active and latent failures during the Nanticoke incident. This failure might have been prevented by barriers that were intended to avoid the release of fuel or by visual inspections once the initial fire had started. The first of these barriers was compromised by the engineer's active failure to ensure a fuel-tight seal for the filter gasket when he restarted the engine. The visual inspections were jeopardised by the restricted field of view that was afforded by the Chief Engineer's and Mechanical Assistant's decision to remain in the Engine Control Room.

In order for an active failures to occur it is necessary for a number of preconditions to be satisfied. These preconditions, typically, relate to the general failure types that were introduced in previous sections. Figure 11.12 provides several examples of this aspect of Tripod-Beta modelling. The watchkeeping engineer's difficulties in achieving a fuel-tight seal were exacerbated by the lack of new, spare copper washer gaskets. This precondition stemmed from a latent failure to identify the importance of these items within the spare parts inventory . This latent failure can, in turn, be associated with the *hardware* general failure type. These hardware failures stem 'from the provision of inadequate equipment and materials' and are the result 'of poor stock control, of problems in the supply chain, of component defects etc'. It can also be argued that the failure to ensure an adequate stock inventory helped to create and was created by *error enforcing conditions*. The fact that the engineer was forced to anneal an existing gasket introduced additional sub-tasks into the preventive maintenance programme. It can argued that this reduced the amount of time available for monitoring and inspection of the generator after it had been reassembled.

A number of further preconditions contributed to the engineer's decision to restart the generator, in spite of the problems that they subsequently reported for their maintenance activities. The modifications to the fuel cover removed the seating groove that helped to ensure an adequate seal. The Watchkeeping Engineer also failed to find any leaks during their initial observation of the generator after the preventive maintenance had been completed. These preconditions are, in turn, be associated with underlying general failure types. Unlike the problems with the stock inventory, mentioned above, it is possible to identify a number of common failure types that may have affected both of these preconditions . For example, *housekeeping* failures relate to problems that have been known for a long time and which have not been adequately addressed. It can be hypothesised that the Engineer did not express concern over the modifications to the forward fuel cover nor did they conduct prolonged inspections of the reassembled generator because the problems that they experienced in obtaining a seal were not unusual. A similar argument might also justify the use of the *communication* and *training* general failure types to characterise the reasons why key personnel failed to report the problems that they faced during maintenance procedures.

The previous paragraphs illustrate the way in which an informal argument must be constructed to explain and justify the decisions and judgements that are represented in Figure 11.12. This is important if other analysts are to understand and accept the reasons why, for instance, the fail-

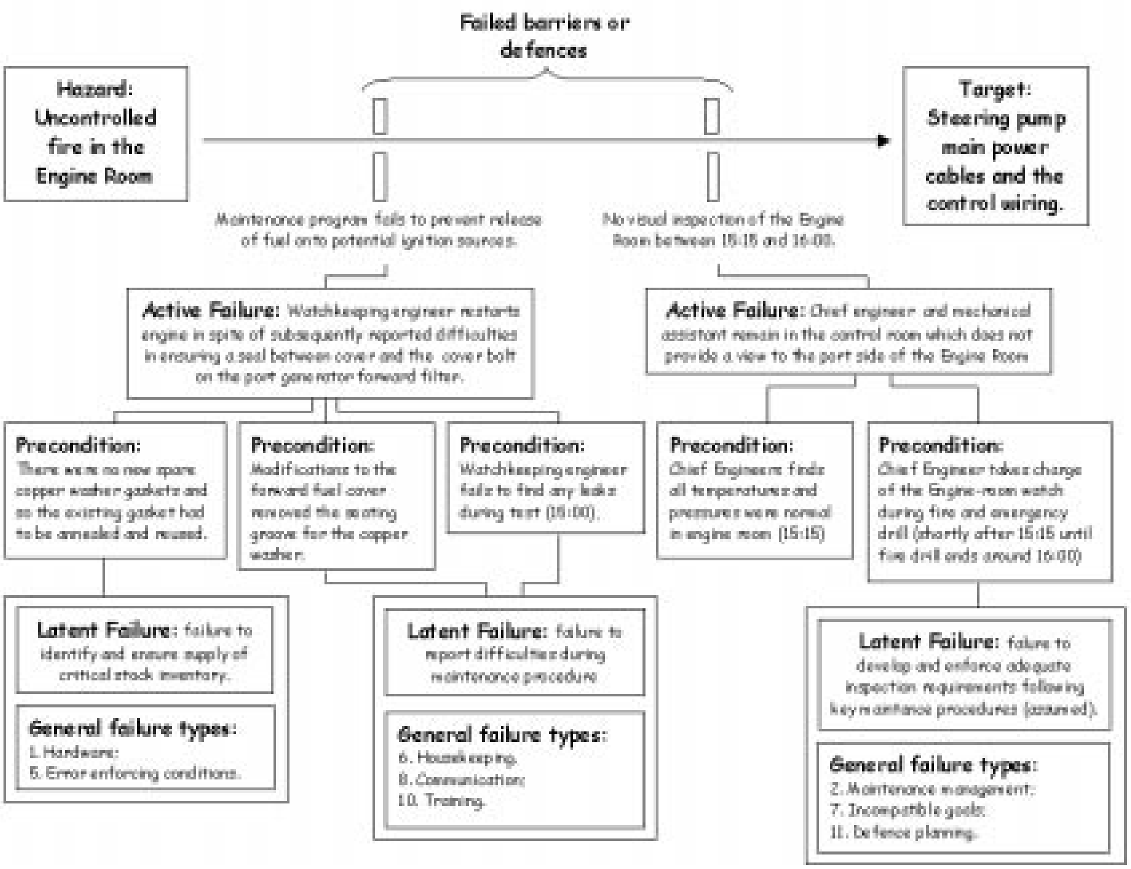
Failed barriers or defences
Hazard: Uncontrolled fire in the Engine Room
Target: Steering pump main power cables and the control wiring.
Maintenance program fails to prevent release of fuel onto potential ignition sources.
No visual inspection of the Engine Room between 15:15 and 16:00.
Active Failure: Watchkeeping engineer restarts engine in spite of subsequently reported difficulties in ensuring a seal between cover and the cover bolt on the port generator forward filter.
Active Failure: Chief engineer and mechanical assistant remain in the control room which does not provide a view to the port side of the Engine Room
Precondition: There were no new spare copper washer gaskets and so the existing gasket had to be annealed and reused.
Precondition: Modifications to the forward fuel cover removed the seating groove for the copper washer.
Precondition: Watchkeeping engineer fails to find any leaks during test (15:00).
Precondition: Chief Engineers finds all temperatures and pressures were normal in engine room (15:15)
Precondition: Chief Engineer takes charge of the Engine-room watch during fire and emergency drill (shortly after 15:15 until fire drill ends around 16:00)
Latent Failure: failure to identify and ensure supply of critical stock inventory.
General failure types:
1. Hardware;
5. Error enforcing conditions.
Latent Failure: failure to report difficulties during maintenance procedure
General failure types:
6. Housekeeping.
8. Communication;
10. Training.
Latent Failure: failure to develop and enforce adequate inspection requirements following key maintance procedures (assumed).
General failure types:
2. Maintenance management;
7. Incompatible goals;
11. Defence planning.

Figure 11.12: Tripod-Beta Event Analysis of the Nanticoke Incident (1)

ure to report maintenance difficulties can be seen as an instance of a more general housekeeping problem. Chapter 9 has explained how the need to provide such rationale is a more general requirement for many analytical techniques. It is especially important when these techniques capture subjective judgements about the underlying causes of a mishap, such as the Nanticoke incident. Other investigators may disagree with the allocation of general failure types represented in Figure 11.12. The provision of a free-text rationale for that allocation can, therefore, be used during subsequent analysis. Additional evidence may also be sought to support assertions about the state of the seating groove during previous maintenance procedures and about the more general reporting of maintenance problems onboard the Nanticoke before this incident.

Preconditions can be thought of as causal factors . They are necessary for an active failure to occur. For instance, the lack of any spare, new gaskets was a necessary precondition for the Watchkeeping Engineer's failure to ensure a seal. Individual precondition need not, however, provide sufficient conditions for an active failure to occur. For example, if the Engineer had detected a leak during their subsequent tests then they might not have decided to restart the generator even if they had been forced to re-use an annealed gasket. The necessary and sufficient conditions for an active failure are represented by the conjunction of all of the preconditions associated with that failure. For example, it was necessary for there not to be any spare gaskets and for modifications to have removed the seating groove and for tests to indicate there were no leaks in order for the Watchkeeping engineer to start the generator. This analysis suggests further links between Tripod-Beta and other techniques, such as ECF analysis and MES, that exploit counterfactual reasoning. For each precondition, analysts must be sure that the associated active failure would not have occurred if that precondition had not been satisfied.

Figure 11.12 represents preconditions that can explain the Chief Engineer's and the Mechanical Assistant's failure to monitor the port side of the Engine Room. The Chief Engineer observed normal temperature and pressure readings in the Engine Room at 15:15. As can be seen, this precondition is not associated with a latent failure or with a general failure type . This is justified because it does not represent a failure. The Chief Engineer correctly monitored the available readings. This was as a precondition for the active failure because it may have reassured him that there were no problems after the preventive maintenance had been completed around 15:00.

The failure to monitor the port side of the engine room may also have been caused by the change in watch that occurred during emergency and fire drills. It is normal practice for Chief Engineers on merchant ships to assume control of the Engine Room during fire and emergency drills. This enables other members of the crew to participate in the exercise while ensuring that normal watchkeeping activities are not compromised. In the Nanticoke incident, the Chief Engineer relieved the watchkeeping engineer who had completed the generator maintenance. This enabled the watchkeeping engineer to proceed to his fire station. This hand-over may, however, have played an important role in the development of the fire. It can be argued that the fire and emergency drills created a context in which the crew were less likely to perform their normal inspection activities. Such interruptions to normal operating procedures can often result in reduced vigilance. Fire and emergency drills provide opportunities for social interaction that are less frequent under the demands of everyday operation. It can also be difficult to ensure that adequate information is handed over from one operator to another. In particular, the Watchkeeping Engineer did not report their difficulty in obtaining a fuel-tight seal. If these concerns had been expressed then the Chief Engineer might have maintained a direct visual observation of the Port-side generator. All of these concerns might have been addressed by the use of operating procedures to ensure that an adequate watch was maintained during the fire and emergency drill [621]. As before, this latent failure is associated with a number of more general failure types. It reflects a failure in *maintenance management*, a problem with *incompatible goals* and potentially with *defence planning*. The maintenance management concerns centre on the need to specify and follow adequate monitoring guidelines during the fire drill after the generator had been restarted. Mishaps are likely to occur if individuals are preoccupied or have goals that conflict with those that are intended to ensure the safety of the system in which they operate. It can be argued that the Chief Engineer's role in assuming the watch during the fire drill might have introduced social or technical demands that impaired their ability to continue monitoring the engine room. Finally, incident can also occur if there are 'deficiencies in the detection, mitigation and remedial actions that are taken in the aftermath of an incident'. This general failure type summarises the role that the active failure played in the incident as a whole, the crews' failure to monitor the port side of the engine room delayed the detection of the fire while it was still taking hold.

Tripod-Beta offers a number of benefits for the causal analysis of safety-critical incidents. In particular, the graphical representation of defences helps to ensure that analysts explicitly consider the way in which active and passive failures combine to jeopardise potential barriers. This is important because other techniques, such as ECF and MES, only consider barriers in an indirect manner. It is possible, however, to raise a number of minor caveats about the manner in which defences are represented in Tripod-Beta. Previous applications of this technique focus on the way in which defences have failed. For example, Figure 11.12 shows how the Nanticoke incident stemmed from a failure to prevent the release of fuel onto a potential ignition source and from a failure to inspect the engine while the resulting fire took hold. A continuing theme in this book has, however, been that near-miss incidents also provide important information about successful defences. This is important if engineers and designers are to accurately assess whether or not those defences can be relied upon to ensure the future safety of a potential target. Figure 11.13, therefore, shows how the conventional use of Tripod-Beta can be extended to consider the role of successful defences and barriers as well as those that are known to have failed. In spite of the Chief Engineers failure to perform a direct visual inspection of the port side of the engine room between 15:15 and 16:00, he did notice the high cooling temperature water alarm that eventually promoted the crews' response to this incident. Figure 11.13 also illustrates a number of additional defences that were not tested during this incident . This is important because, as we shall see, any recommendations must also consider what

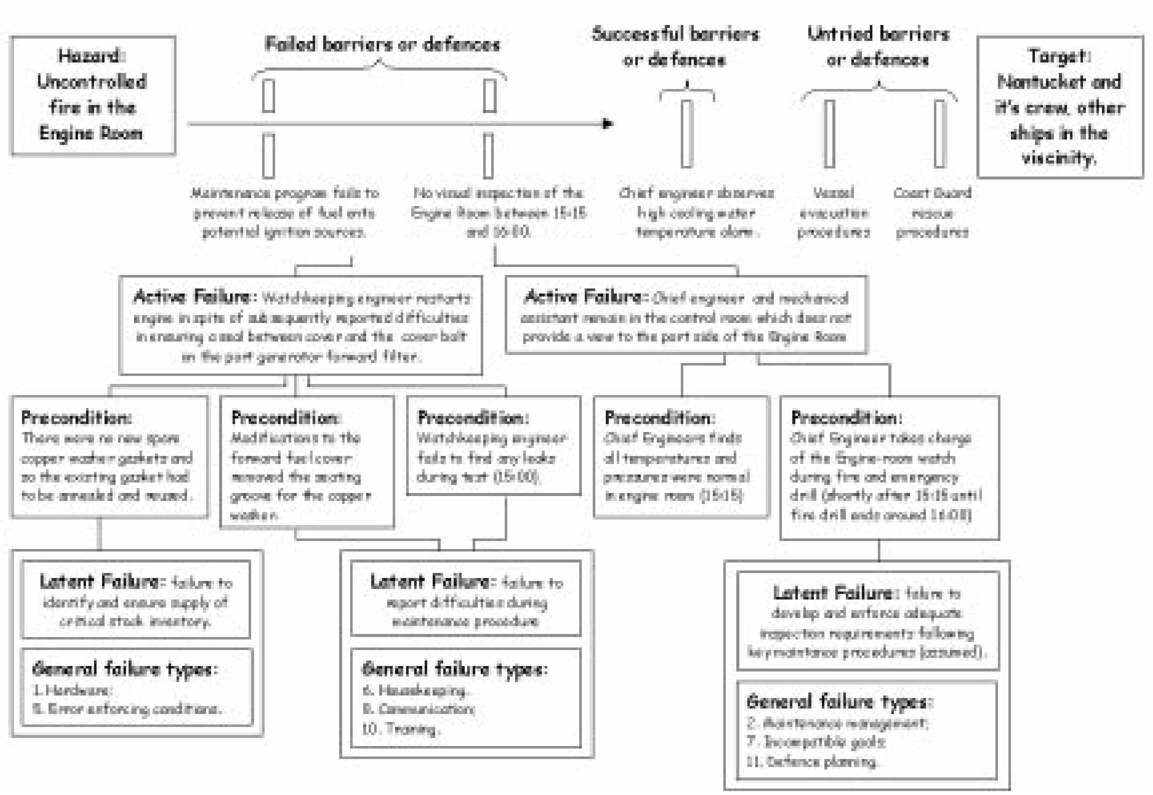
Hazard: Uncontrolled fire in the Engine Room
Failed barriers or defences
Successful barriers or defences
Untried barriers or defences
Target: Nantucket and it's crew, other ships in the viscinity.
Maintenance program fails to prevent release of fuel onto potential ignition sources.
No visual inspection of the Engine Room between 15:15 and 16:00.
Chief engineer observes high cooling water temperature alarm.
Vessel evacuation procedures
Coast Guard rescue procedures
Active Failure: Watchkeeping engineer restarts engine in spite of subsequently reported difficulties in ensuring a seal between cover and the cover bolt on the port generator forward filter.
Active Failure: Chief engineer and mechanical assistant remain in the control room which does not provide a view to the port side of the Engine Room
Precondition: There were no new spare copper washer gaskets and so the existing gasket had to be annealed and reused.
Precondition: Modifications to the forward fuel cover removed the seating groove for the copper washer.
Precondition: Watchkeeping engineer fails to find any leaks during test (15:00).
Precondition: Chief Engineers finds all temperatures and pressures were normal in engine room (15:15)
Precondition: Chief Engineer takes charge of the Engine-room watch during fire and emergency drill (shortly after 15:15 until fire drill ends around 16:00)
Latent Failure: failure to identify and ensure supply of critical stock inventory.
General failure types:
1. Hardware;
5. Error enforcing conditions.
Latent Failure: failure to report difficulties during maintenance procedure
General failure types:
6. Housekeeping.
8. Communication;
10. Training.
Latent Failure: failure to develop and enforce adequate inspection requirements following key maintance procedures (assumed).
General failure types:
2. Maintenance management;
7. Incompatible goals;
11. Defence planning.

Figure 11.13: Tripod-Beta Event Analysis of the Nanticoke Incident (2)

might have happened if the successful defence had also been compromised. Tripod-Beta also offers a number of further benefits. These can be summarised as follows:

- *focussed use of a time-line.* The event analysis component of the Tripod-Beta technique includes a time-line that shows some similarity with those that are used in ECF analysis, in MES and STEP. This is represented by the horizontal arrow that potentially connects a hazard to a target in Figures 11.12 and 11.13. Unlike the alternative analysis techniques, Tripod-Beta focuses on those events that are associated with the failure of defences. This considerably simplifies the modelling of an incident or accident. The sparse approach advocated by Tripod-delta also omits information, such as the actors involved in an event, that forms an important component of the MES and STEP approaches;

- *explicit representation of active and latent failures.* The distinction between latent and active failures reflects much recent research into the nature of technological failure [701, 362]. It is intended to move the focus of an analysis away from the individual failures that characterise a particular incident to look for more general managerial and organisational causes. In other techniques, such as ECF analysis, this distinction is only recognised through the use of auxiliary techniques such as Tier analysis;

- *support for a checklist approach to root cause analysis from the eleven general failure types.* The general failure types in Tripod-Beta are similar to the leaf nodes within PRISMA classification models. They describe a number of recurring 'root causes'. They support analysts by directing their attention to recognised causes of previous incidents. This, in turn, can encourage greater consistency between investigators than might otherwise be possible with techniques that do not exploit a checklist approach.

- *balance between a high level of abstraction in the general failure types and more specific information from the use of preconditions.* It is possible to contrast the eleven general failure types supported by Tripod-Beta with the one thousand five hundred items in a full MORT diagram. It is far easier to perform an initial analysis using this limited number of general failure types than it is to perform an exhaustive search through a MORT diagram. Conversely, it can be more difficult to identify general failure types that accurately characterise the particular root causes of incidents within a particular industry or organisation. PRISMA avoids this problem by combining a relatively simple checklist, which is similar to aspects of Tripod-Beta, with a recommendation that analysts extend the classification scheme to reflect local conditions within particular industries. For instance, Van Vuuren's medical checklist includes an item for 'patient related factors' [844]. This item is not included within Van der Schaaf's PRISMA taxonomy for chemical incidents [841]. It can, however, be argued that such differences can introduce important inconsistencies between the results of causal analyses that were obtained using different classification schemes. Tripod-Delta avoids some of these problems by explicitly representing the relationship between specific details of an incident, in the annotations associated with active failures and with preconditions, and the more general root causes. These annotations can be used to stress particular aspects of an incident that cannot easily be captured using the restricted palette offered by the eleven general failure types.

- *tool support.* Finally, the application of Tripod-Beta is supported by a number of computer-based tools. This is significant because these systems can also be integrated with the constructive use of Tripod-Delta as part of a wider safety management programme. The Tripod-Beta tools provide a number of internal consistency checks that help analysts to construct the event analysis diagrams, illustrated in Figures 11.12 and 11.13. It is important to stress, however, that our analysis was conducted without the use of these tools. This provided greater flexibility, for example in the representation of successful barriers in Figure 11.13, that might not be so desirable if an organisation were keen to ensure greater consistency between the event diagrams that were produced by incident analysts.

The benefits of Tripod-Beta analysis must be balanced against a number of potential problems. In particular, this technique raises concerns that are similar to those that motivated Benner to omit conditions from the STEP approach. It can be difficult to distinguish between active failures and preconditions. For instance, Figure 11.12 argued that the Watchkeeping Engineer's failure to find any leaks was a precondition for their active failure in restarting the engine without reporting a potential maintenance problem. It might be argued that the failure to detect any leaks should be classified as an active failure in its own right. This would result in a graph in which an active failure is the result of both preconditions and of active failures. Each of these active failures might, in turn, be the result of further preconditions and further active failures and so on. The ECF analysis in Chapter 10 has illustrated the complexity of a similar approach. This technique might have even worse consequences for Tripod-Beta; ECF charts do not distinguish between active and passive failures.

To summarise, preconditions introduce a potential ambiguity into Tripod-Beta modelling. They capture information about the state of a system; modifications to the forward fuel cover removed the seating groove for the copper washer. They also capture event-based information; watchkeeping engineer fails to find any leaks during test (15:00). This creates ambiguity because these events may themselves represent active failures that can be associated with further pre-conditions. In practice, it is possible to develop a number of heuristics that reduce the consequences of such ambiguity. For instance, Figures 11.12 and 11.13 only consider the preconditions of those active failures that are directly associated with the failure of particular barriers. The analysis does not consider the preconditions of a precondition. If analysts wanted to consider the Watchkeeping Engineer's failure to find any leaks then that event would have to be associated with the failure of a particular barrier at the top level of the Tripod-Delta diagram.

The application of Tripod-Beta has also shown how analysts must provide considerable additional documentation to support the diagrammatic form illustrated in Figures 11.12 and 11.13. In particular, it is important to explain why particular latent failures can be associated with general failure types. Similarly, rationale must be provided so that other analysts can understand the relationship between a latent failure and a particular precondition. Our analysis of the Nanticoke case study illustrated this issue when we considered the possible impact that the fire and emergency drills might have had upon the monitoring of the port side of the engine room. In order to interpret the relationship between the precondition, latent failure and general failure types, analysts must understand the manner in which responsibilities and tasks are routinely handed-over so that other members of the crew can participate in the drill. It was also necessary to draw upon evidence from previous failures to explain the problems that can arise from the transfer of information during such hand-overs. This additional information illustrates the manner in which the Tripod-Beta event analysis diagram is not an end in itself. It provides a high-level framework for the causal analysis of incidents and accidents. It does not, however, replace the more general inferential and reasoning skills that are established by expertise and training in incident analysis.

11.3 Mathematical Models of Causation

Previous sections have introduced a number of semi-formal techniques that are intended to support the causal analysis of safety critical incidents. They can be classified as 'semi-formal' because it can be difficult to develop a coherent set of rules to describe the syntax and semantics of the associated notations. For instance, we have identified some of the problems that can arise when attempting to construct a precise definition of the preconditions that form an important component of Tripod's event analysis diagrams. Similarly, it can be difficult to derive a precise definition for what can and what cannot be represented in the leaf nodes of a causal tree. Investigators are free to use natural language annotations. This increases the flexibility of the approach. It can, however, also introduce potential ambiguity and inconsistency if a team of investigators must cooperate in the construction of a shared tree during a PRISMA analysis. A number of organisations have responded to these problems by developing more formal, mathematically based, causal analysis techniques.

11.3.1 Why-Because Analysis (WBA)

Why-Because Analysis stems from an initiative to increase the objectivity of accident investigations by encouraging "rigorous causal analysis" [469]. The technique is based around two complementary stages. These can be summarised as follows:

1. *Construct the Why-Because Graph.* The first stage in the WBA involves the construction of a graph that is intended to capture the significant causal relationships that led to an incident. The causal relationships are identified using the counterfactual reasoning that has been a feature of previous approaches. The method is, however, supported by a formal semantics for causation that is based on that provided by the philosopher and logician David Lewis, mentioned in previous chapters [490, 491].

2. *Prove that the Graph is Sufficient and Correct.* The previous techniques that have been presented in this chapter and in Chapter 10 would stop after stage 1 of the WBA. In contrast, however, this logic-based technique provides procedures for ensuring that the causal relations in a Why-Because graph actually satisfy the semantics for causation that is implied by Lewis' underlying model. In other words, there are rules for showing that the model of an incident reflects Lewis' view of causation. These techniques can also be used to ensure that there is a sufficient causal explanation for each identified fact that is not itself a root cause [469].

The following pages provide a brief introduction to these two stages of analysis. It is important to emphasise, however, that the benefits provided by the mathematical underpinning of WBA can also important impose considerable upon the analyst. The various stages of the technique can appear to be extremely complex even for investigators who have a background in mathematical logic. As we shall see, therefore, this approach may be most suitable for near-miss incidents that might under other circumstances have resulted in high-consequence accidents.

As mentioned, the first stage of WBA involves the construction of a graph that is intended to capture the causal relationships that lead to incidents and accidents. The nodes of these graphs represent four different factors: states; events; processes and non-events [499]. States are represented by collections of state predicates. These can be thought of as sentences that are true in that state. For example, the ignition of the Nanticoke fire might be represented by state in which it was true that 'fuel is being sprayed under pressure from the forward fuel filter of the port generator'. WBA uses angled brackets to denote individual states, $\langle \mathit{State} \rangle$. Events represent changes in state. For instance, the deployment of the Halon system is an event that transformed the state of the Nanticoke from one in which there was a fire to one in which there was no fire. WBA uses brackets to denote individual events, $[\mathit{Event}]$. Processes can be defined to describe mixtures of states and events that have some bounded duration. For example, the Nanticoke incident can be described in terms of a process in which the maintenance event transformed the state of the system into one in which a fire could occur. The ignition event changed the state of the system into one in which a fire was taking place and so on. WBA uses curling brackets to denote processes, $\{\mathit{Process}\}$. Finally, as we have seen, it is often necessary to consider the impact that errors of omission have upon the course of an incident. WBA, therefore, provides non-events using the following notation $(non - events)$

WBA proceeds by developing a history of the incident. Successive states of the system are liked using a temporal ordering relation that is denoted using the $\hookrightarrow$ symbol. For more information on the semantics of the $\hookrightarrow$ operator, see Lamport [473]. For now it is sufficient to observe that it forms part of a more complex Explanatory Logic that was developed by Ladkin and Loer to provide means of formally demonstrating the correctness of a causal argument [470, 499]. The initial stages of the Nanticoke case study can be represented by the following high-level history:

$$\langle \mathit{Maintenance} \rangle \hookrightarrow \langle \mathit{Fire} \rangle \hookrightarrow [\mathit{Deploy\ Halon\ System}] \tag{11.1}$$

It is important to emphasise that the temporal ordering, captured by the $\hookrightarrow$ symbol, does not represent causality. Loer illustrates the distinction between causation and temporal sequence [499]. A traffic-jam may occur immediately after I leave the highway, however, there need not be any causal relationship between these two events unless I have parked my car across the carriage-way. A number of axioms can be used to describe important properties that must exist between temporal and causal relations. For example, if a causal chain exists such that A causes B then the first element of this causal chain, A, must occur before the last element, B. This leads to the following inference rule:

$$\frac{A \Rightarrow\!\!\succ^* B}{A \hookrightarrow B} \tag{11.2}$$

If we know that A causes B, $A \Rightarrow\!\!\succ^* B$, then we can also conclude that A must precede B, $A \hookrightarrow B$. If this rule were not to hold then past events could be the result of situations that still lie in the future!

To summarise, we would like to be able to construct a causal model of an incident using the $\Rightarrow\!\!\succ^*$ operator. Most primary and secondary investigations result in temporal models, similar to those proposed in Chapter 9. These describe sequences that can be represent using the $\hookrightarrow$ operator. Unfortunately, there is no automatic means of translating temporal sequences into causal relations. Many different causal chains can produce the same high-level temporal sequence. For instance, the (11.1) sequence might have been caused by maintenance to the starboard generator, to the transmission system and so on. Investigators must apply their skill and expertise to identify the causes of the temporal sequences that can be reconstructed in the aftermath of an incident. Fortunately, WBA provides an informal procedure that helps in this task. This process starts by asking *Why did the final event in the sequence occur?*. For the Nanticoke example in (11.1) this would yield:

Why was the Halon system deployed?.
Because the second fire party withdrew from fighting the fire.

The analysis continues by asking, in turn, why did the second fire party withdraw? This was because they were ordered to abandon their attempt to extinguish the fire. As mentioned, the key

Why-Because questions are intended to guide the process by which the temporal $\hookrightarrow$ sequences are translated into more detailed causal relations, $\Rightarrow\!\!\!\succ^*$. However, this process may also help to identify factors that were not considered during the initial temporal sequence. For instance, the previous questions helped to identify that the failure of the second fire party was a reason why the Halon system was deployed. Our previous analysis did not include any information about either the first or the second fire party. Figure 11.14 illustrates how this recursive analysis can be used to identify the reasons why the Halon system was deployed. The first fire party's attempt to use carbon-dioxide extinguishers was beaten back by the heat of the fire. This led to a second fire party attempting to use charged hoses. This attempt was ordered out of the engine room which then led to the Chief Engineer discharging the Halon system.

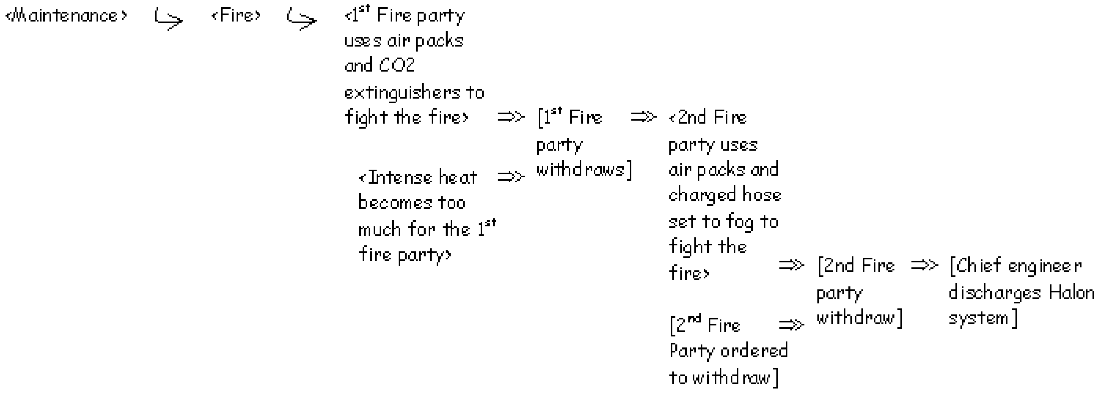
⟨Maintenance⟩
⟨Fire⟩
⟨1st Fire party uses air packs and CO2 extinguishers to fight the fire⟩
⟨Intense heat becomes too much for the 1st fire party⟩
[1st Fire party withdraws]
⟨2nd Fire party uses air packs and charged hose set to fog to fight the fire⟩
[2nd Fire Party ordered to withdraw]
[2nd Fire party withdraw]
[Chief engineer discharges Halon system]

Figure 11.14: Why-Because Graph Showing Halon Discharge

A number of observations can be made about the Why-Because graph illustrated in Figure 11.14. As can be seen, the maintenance and fire states that were identified in (11.1) continue to be connected by the sequence relation, $\hookrightarrow$. However, the causal analysis has helped to identify states and events that are causally related, $\Rightarrow\!\!\!\succ$. Formally, $\Rightarrow\!\!\!\succ$ is the transitive closure of $\Rightarrow\!\!\!\succ^*$. Informally, $A \Rightarrow\!\!\!\succ B$ denotes that A is a direct causal factor of B. $A \Rightarrow\!\!\!\succ^* B$ represent situations in which there may be intermediate or 'knock-on' causal relations. For example, in Figure 11.14 we can say that the withdrawal of the first fire party is a direct causal factor behind the 2nd fire party's use of the hoses to fight the fire, denoted using $\Rightarrow\!\!\!\succ$. In contrast, the withdrawal of the first fire party is a knock-on cause of the Chief Engineer's action to discharge the Halon system, denoted using $\Rightarrow\!\!\!\succ^*$.

Why-Because graphs, typically, use a numerical indexing system rather than the free-text labels that are shown in Figure 11.14. $\langle Maintenance \rangle$ might be denoted by $\langle 1 \rangle$, $\langle Fire \rangle$ by $\langle 2 \rangle$ and so on. This has not been done because the graph is relatively simple and the labels are intended to help the reader trace the causes of the Halon deployment. However, this approach quickly becomes intractable as the scope of the graph increases.

It is possible to perform a number of consistency checks using the formal rules that underpin the graphical notation that is provided by Why-Because graphs. The simplest of these involves checking that the causal relationships are consistent with the previous temporal order described in (11.1) using the $\hookrightarrow$ operator. Or more formally, the analyst must ensure that the transitive closure of the causal relations in Figure 11.14 continue to preserve the temporal sequence of (11.1) [499].

It should also be noted that, as might be expected, it can be difficult to determine how best to represent an incident using the four factors that form the nodes of a Why-Because graph: $\langle State \rangle$; $[Event]$; $\{Process\}$; $(Non-events)$. As mentioned, analysts must decide whether a particular aspect of an incident is best represented as a state, en event, a process or as a non-event. It is relatively straightforward to distinguish an event from a non-event. It can, however, be more complex to determine what is an event and what is a process. For example, Figure 11.14 shows that the discharge of the Halon system was a discrete event. It can also be argued that the task of deploying this form of extinguisher is more likely to have been composed from a sequence of events and could, therefore, be better represented as a process. The Chief Engineer must form the intention to deploy the system. He must then ensure that everyone is accounted for and that none is left in the area in which the system will be deployed. There may have been a confirmation protocol to inform the Captain the system was to be deployed etc. This decision between an event or a process is typical of the choices that must be made when using many different causal analysis techniques. It reflects the level of detail that the analyst considers to be necessary when constructing a model of an incident or accident. The key point is that the model explicitly represents this information so that other analysts can review their colleague's view of an incident and, if necessary, request additional detail.

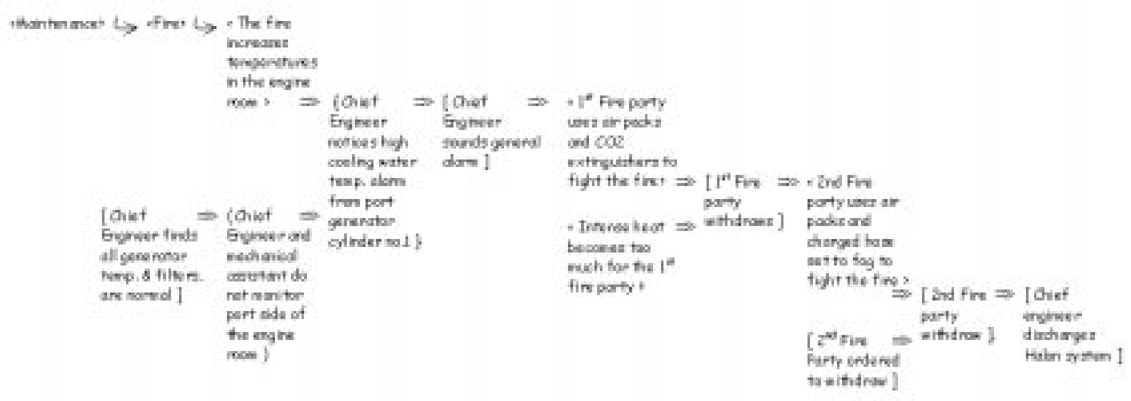

<Maintenance>
<Fire>
< The fire increases temperatures in the engine room >
{ Chief Engineer notices high cooling water temp. alarm from port generator cylinder no.1 }
[Chief Engineer sounds general alarm]
< 1st Fire party uses air packs and CO2 extinguishers to fight the fire >
[1st Fire party withdraws]
< Intense heat becomes too much for the 1st fire party >
< 2nd Fire party uses air packs and changed hose set to fog to fight the fire >
[2nd Fire party withdraw]
[2nd Fire Party ordered to withdraw]
[Chief engineer discharges Halon system]
[Chief Engineer finds all generator temp. & filters are normal]
(Chief Engineer and mechanical assistant do not monitor port side of the engine room)

Figure 11.15: Why-Because Graph for the Nanticoke Alarm

Figure 11.15 extends the analysis to consider the reasons why the first fire party was called on to combat the fire in the first place. As can be seen, they were responding to a general alarm. Why had the general alarm been issued? Because the Chief Engineer had noticed the high cooling water temperature alarm. Why had the Chief Engineer noticed this alarm? Because the fire had increased temperatures in the engine room. The Chief Engineer had also noticed this alarm because he and the mechanical assistant had not monitored the port side of the engine room and so had not noticed the fire before it took hold. Why had they not monitored the port side of the engine room? Because an initial inspection had not shown anything unusual with the generators.

Figure 11.15 illustrates a number of further properties of Why-Because graphs. For instance, the reason that the Chief Engineer eventually observes a high cooling water temperature alarm is because they and the Mechanical Assistant fail to monitor the port side of the engine room. This is denoted as a $(non - event)$. In order to capture the semantics of these non-events, the Explanatory Logic of WBA draws upon deontic arguments of obligation and permission. The crewmembers violated the procedures and norms that obliged them to maintain a visual watch over the engine room. Ladkin and Loer provide a full justification for this application of deontics [470]. For now it is sufficient to observe that WBA provides a meta-rule that is intended to guide investigators in the identification of these non-events. Investigators must explicitly add a non-event, (E), to the history of states if $O\langle E\rangle$ is derivable and E does not occur, where O represents deontic obligation [470]. Figure 11.15, therefore, includes the non-event Chief Engineer and Mechanical Assistant do not monitor port side of the engine room.

Figure 11.15 also illustrates the way in which the $\{Process\}$ format provides powerful abstractions that can be used to describe complex causal sequences. For instance, there are likely to be a number of perceptual and cognitive mechanisms that led the Chief Engineer to notice the high cooling water temperature alarm. Subsequent analysis could recruit human factors experts to identify these factors. During any initial analysis, however, the details of this cognitive and perceptual process can be denoted as {Chief Engineer notices high cooling water temperature alarm from port generator cylinder number 1}. The process form is also used to represent the human factors mechanisms that led the Chief Engineer to conclude that there were no problems during their initial inspection of the engine room.

Figure 11.16 illustrates the results of applying WBA to the factors in the temporal sequence that was introduced in (11.1). As mentioned, the formal underpinnings of this analytical technique are intended to ensure that investigators can represent and reason about the products of their investigations. This helps to ensure that errors are avoided during the construction of relatively complex Why-Because graphs, such as that illustrated in Figure 11.16. These 'quality control' procedures take two principle forms. The first approach uses information about each node to ensure that a Why-Because graphs satisfy a number of high level properties. For example, investigators must ensure that each node has at least one causal factor that represents an event. They must also ensure that each node is classified exclusively as one of the four factors mentioned above.

Additional constraints can be imposed, for example, to ensure that investigators minimise the use of processes wherever possible. This injunction is justified because processes should not be used as a 'catch all' when investigators find it difficult to discriminate between events and states. In other words, they should not be used to mask or hide aspects of an incident that ought to be the subject of a more detailed investigation. For instance, Figure 11.16 might be refined to consider what exactly attracted the Chief Engineer's attention to the cooling water high-temperature alarm. It is important to stress that WBA was developed to support the investigation of accidents rather than near-miss incidents. More limited analytical and investigatory resources may, therefore, prevent individuals from obtaining the evidence that is necessary to resolve processes into their component states and events. There may be other processes, such as { 1st party decides to withdraw } in Figure 11.16, that may involve complex perceptual, cognitive and physiological 'states' or 'events'. Such processes are difficult to analyse. As we have seen, investigators may be forced to assume intention from observed behaviour. The proponents of WBA have developed the Perception, Attention, Reasoning, Decision, Intention and Action (PARDIA) model to help analyse such processes. Loer stresses that PARDIA should be used to classify rather than to understand error [499]. This is a fine distinction given that he constructs a normative model of intention. The details of this model are beyond the scope of the current work. It should be noted, however, that PARDIA focuses on cognitive, perceptual and physiological attributes of single operators. It, therefore, suffers from some limitations when applied to team-based incidents and accidents. As we have seen, group dynamics often lead to situations in which team-based behaviour cannot simply be described as the 'sum of its members'.

Automated tools have been developed to assist with the checks, described above. This is important because an error in writing an event node as a state, or an event node which only has states as causal factors, can result in a consistency review on the sub-graph leading to this node. Formal proof techniques provide an alternative means of ensuring the integrity of WBA. As we shall see, however, the costs of performing this analysis may dissuade investigators from going 'the full distance' on this form of analysis [469].

A number of benefits can be derived from the close relationship between WBA and philosophical work on the nature of causation [499]. For instance, investigators must often explain why one version of events is more plausible than another. Lewis has proposed the idea of *contrastive explanation* as a technique that can be used to support these arguments about plausibility [492]. If we have to decide between two versions of an incident we must assess the evidence that is derived through a primary and secondary investigations. In addition, we can also contrast the causal explanation of those histories as revealed using techniques such as the WBA. This approach resembles earlier arguments in this chapter; causal analysis often identifies the need to provide further evidence in support of hypothesised causal relations. An important application of this idea is that any causal analysis must not only explain why an incident occurred in a particular way, it must also explain why the system did not function in the manner intended. For example, Loer analyses an incident in which a DC10 landed at Brussels rather than its intended destination of Frankfurt Airport [499]. He uses WBA to contrast the actual incident, in which the aircraft landed in Brussels, with the "deontically-correct" world in which the aircraft was supposed to land at Frankfurt. His analysis proceeds by identifying the earliest contrast between what actually did happen and what was supposed to happen. In this case, the aircraft was transferred to Brussels Air Traffic Control rather than Maastricht . Figure 11.17 shows how this technique might be applied to the Nanticoke incident.

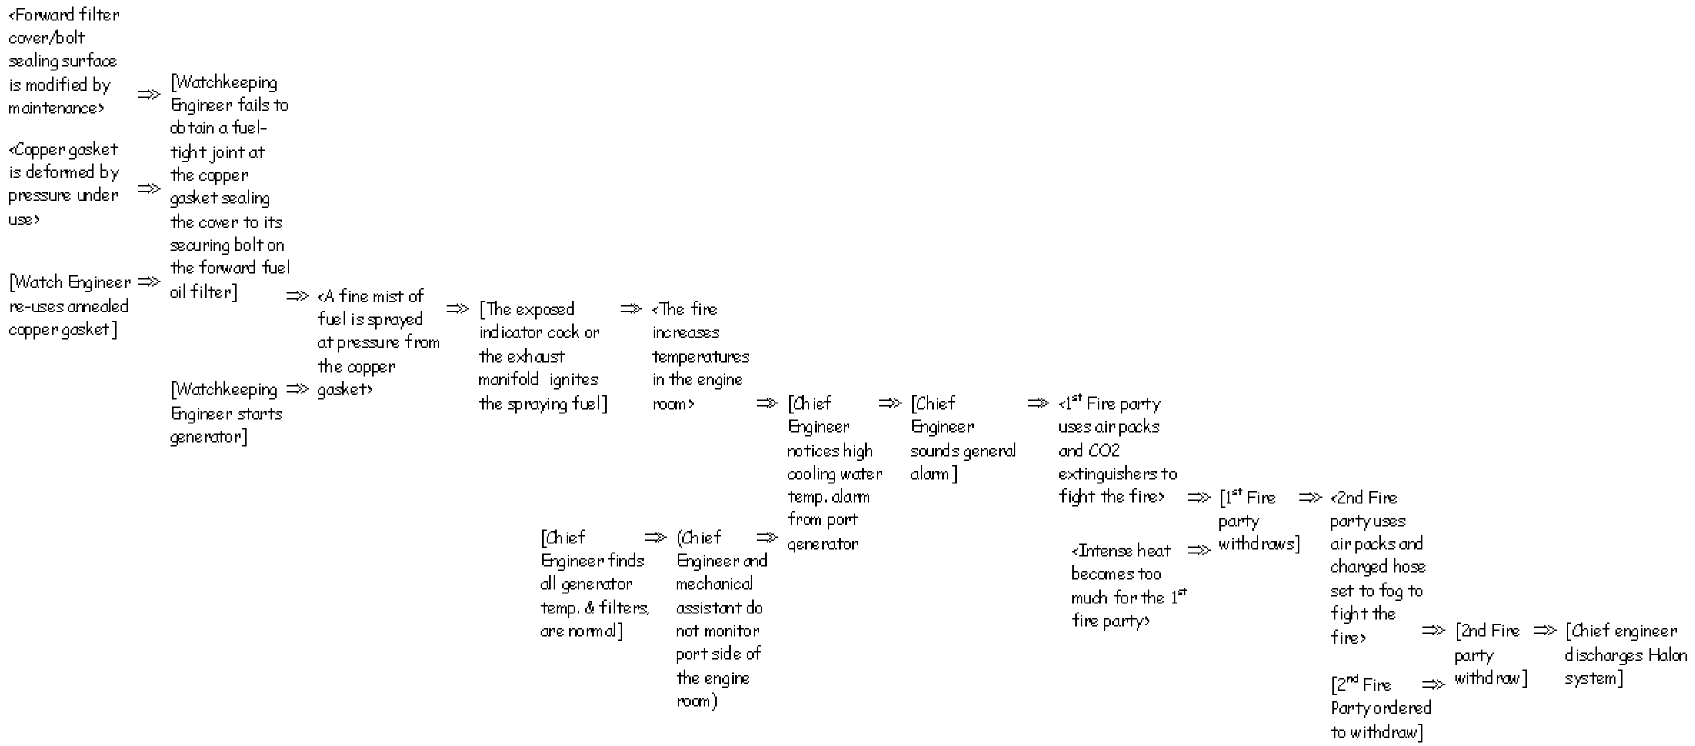
<Forward filter cover/bolt sealing surface is modified by maintenance>
<Copper gasket is deformed by pressure under use>
[Watch Engineer re-uses annealed copper gasket]
[Watchkeeping Engineer fails to obtain a fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter]
[Watchkeeping Engineer starts generator]
<A fine mist of fuel is sprayed at pressure from the copper gasket>
[The exposed indicator cock or the exhaust manifold ignites the spraying fuel]
<The fire increases temperatures in the engine room>
[Chief Engineer finds all generator temp. & filters, are normal]
(Chief Engineer and mechanical assistant do not monitor port side of the engine room)
[Chief Engineer notices high cooling water temp. alarm from port generator
[Chief Engineer sounds general alarm]
<1st Fire party uses air packs and CO2 extinguishers to fight the fire>
<Intense heat becomes too much for the 1st fire party>
[1st Fire party withdraws]
<2nd Fire party uses air packs and charged hose set to fog to fight the fire>
[2nd Fire Party ordered to withdraw]
[2nd Fire party withdraw]
[Chief engineer discharges Halon system]

Figure 11.16: Overview of the Why-Because Graph for the Nanticoke Incident

The maintenance modifications that damaged the forward filter cover/bolt sealing surface occurred before the copper gasket was deformed or the maintenance engineer annealed the gasket. This event might, therefore, provide a good starting point for any contrast between what did happen and what ought to have happened. It can be argued that maintenance personnel should have noticed the damage and reported it through a management system. This should have resulted in the surface being made good before a fire could occur. The ? ⇒? symbols are used to distinguish causal links from this "possible" world in Figure 11.17. In practical terms, this analytical technique is useful because it can be used to identify non-events that might otherwise be omitted from an analysis. We could redraft the graph in Figure 11.17 by replacing the possible worlds with (Maintenance personnel do not notice the damage to the sealing surface). This particular application of contrastive explanations has much in common with barrier analysis. It can be used to explain the failure of a defensive mechanism that was intended to ensure that the system returned to a 'normative' state.

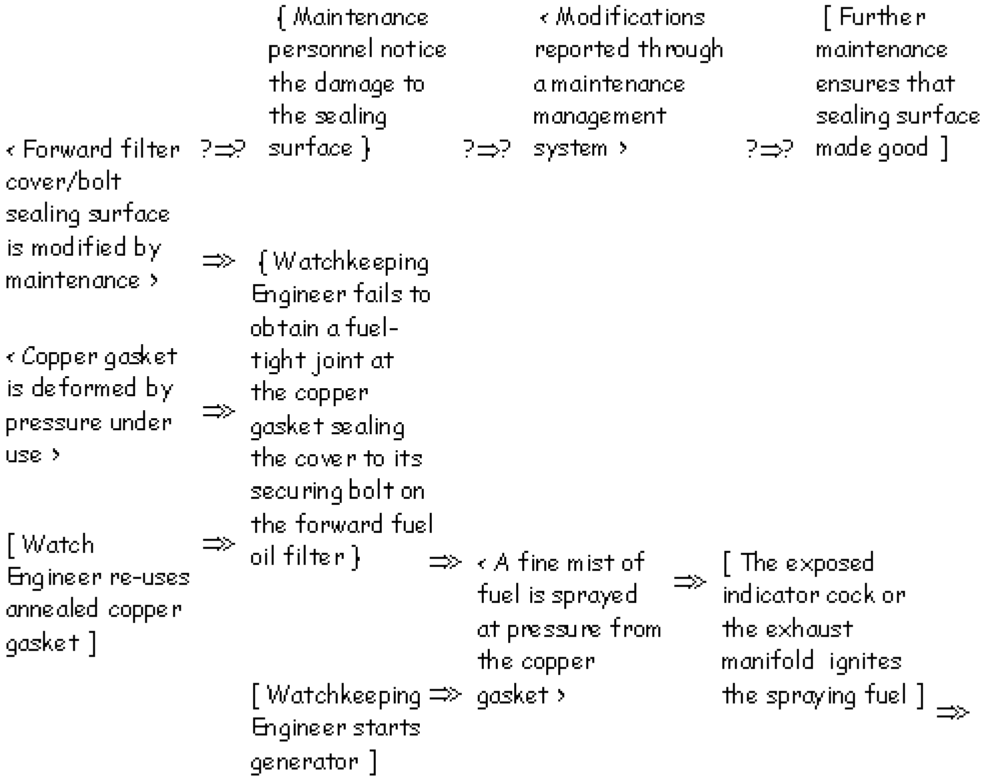
{ Maintenance personnel notice the damage to the sealing surface }
< Modifications reported through a maintenance management system >
[Further maintenance ensures that sealing surface is made good]
< Forward filter cover/bolt sealing surface is modified by maintenance >
?⇒?
?⇒?
?⇒?
⇒
{ Watchkeeping Engineer fails to obtain a fuel-tight joint at the copper gasket sealing the cover to its securing bolt on the forward fuel oil filter }
< Copper gasket is deformed by pressure under use >
⇒
[Watch Engineer re-uses annealed copper gasket]
⇒
⇒
< A fine mist of fuel is sprayed at pressure from the copper gasket >
⇒
[The exposed indicator cock or the exhaust manifold ignites the spraying fuel]
⇒
[Watchkeeping Engineer starts generator]
⇒

Figure 11.17: Possible 'Normative' Worlds for the Nanticoke Incident

A number of minor issues complicate our application of contrastive explanations. Loer's example of this technique is relatively straightforward [499]. The normative and non-normative paths diverge from the point at which the DC10 was handed to Brussels and not Maastricht Air Traffic Control. The Nanticoke incident is not quite so straightforward. It is also important to emphasise that our Why-Because graph ends with the deployment of the Halon system. This is justified because it is important to learn about the resolution of adverse incidents as well as the causes of any failure. One consequence of this is that we cannot simply look for the earliest contrast with a possible world in which the Halon was not deployed. We must also ensure that the alternative 'normative' world avoid a fire. As mentioned, there is a relatively simple divergence in Loer's DC10 case study. In contrast, the Nanticoke case study contains several points at which non-normative and normative behaviours can be distinguished. The Chief Engineer was supposed to monitor the engine room.

The two fire parties were supposed to be deployed before the fire required the use of Halon. By focusing on the earliest contrastive explanation, analysts might miss important lessons about other failures that contributed to the course of an incident.

There are further complications. We have identified the earliest contrast between what did and what should have happened as the maintenance on the filter cover and bolt sealing surfaces. Notice, however, that the previous Why-Because graphs did not specified any temporal sequence over this state and the other two initial causal factors that describe the deformation and re-use of the copper gasket. This sequence was inferred from the evidence that was obtained in the aftermath of the incident. It would, however, also be possible to contrast possible worlds in which the deformation of the gasket was monitored or in which the maintenance engineer did not have to re-use an annealed gasket. As before, this analysis can be used to help investigators explain why particular barriers failed to protect the system. For instance, additional nodes might be introduced into the Why-Because graph to indicate that the absence of a maintenance reporting system explains why these factors were not addressed prior to the incident. It is important, however, that investigators recruit evidence to justify their assertions about these potential defences. For instance, it is not clear that the incident would have been avoided even if the maintenance issues had been effectively reported. The absence of necessary parts or delays in maintenance scheduling might still have led to an adverse occurrence.

The previous paragraphs have described how an informal analysis of alternative possible worlds can be used to distinguish 'normative' from 'non-normative' behaviour. This is useful in identifying ways in which barriers, including regulations, working practices and automated systems, failed to prevent an incident from occurring. As we have seen, however, there are also situations in which investigators cannot distinguish between alternative causal explanations. For instance, we do not know whether the Nanticoke fire was ignited by the exposed indicator tap or by the exhaust manifold. We could use the $? \Rightarrow ?$ notation to describe two divergent causal paths. One might indicate that the indicator tap ignited the fire, the other might represent the exhaust manifold as the ignition source. This can create considerable additional complexity as almost half of the graph would be duplicated. Lewis suggests that these alternative explanations should be ranked by experts using some weighting mechanism [492]. If an alternative explanation was considered sufficiently unlikely then it can be omitted from subsequent analysis. There are a number of concerns about whether this is possible either in the general case or in the example of the Nanticoke fire [671, 469]. Loer advocates the retention of these different paths but acknowledges the consequent complexity [499]. We have, therefore, retained the node labelled The exposed indicator tap or the exhaust manifold ignites the spraying fuel. We have not duplicated the rest of the Why-Because graph, however, because the consequences of these two possible worlds are indistinguishable. Neither of these approaches provides a more general solution to this problem and it remains a subject for future research. It should also be noted that non-determinism complicates the application of all causal modelling techniques. This is most clearly seen in the closing sections of Chapter 10 where we recognised the difficulty of using ECF to model alternative causal hypotheses about the loss of the Mars Climate Orbiter and the Polar Lander.

The previous paragraph argued that investigators can construct different Why-Because graphs to represent alternative causal explanations. These alternative explanations can be thought of as 'possible worlds'. For instance, there is one possible world in which the Nanticoke fire was caused by the indicator tap and another in which it was caused by the exhaust manifold. This notion of alternative possible worlds provides WBA with a semantics for the counterfactual arguments that investigators use to identify causal factors. Chapter 7 distinguished causal factors using the argument:

> A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either.

Lewis [492] recasts this in the following manner:

> "A is a causal factor of B, if and only if A and B both occurred and in the nearest possible worlds in which A did not happen neither did B".

Ladkin and Loer formalise this definition as follows:

$$\begin{array}{c} A \wedge B \\ \neg\, A \;\Box\!\!\rightarrow\; \neg\, B \\ \hline A \Rightarrow\!\!\!\succ B \end{array} \tag{11.3}$$

Informally, $A \;\Box\!\!\rightarrow\; B$ captures the notion that B is true in possible worlds that are close to those in which A is true. As can be seen, (11.3 uses this operator to express the counterfactual component of the Lewis definition. As mentioned, Loer and Ladkin provide a more detailed presentation of this application of Lewis' work [499, 470]. The key point, however, is that logic can be used to provide a clear semantics for informal concepts such as 'cause'. Investigators can also use associated proof rules to ensure both the consistency and sufficiency of informal reasoning about the causes of incidents and accidents.

The formal underpinnings of the Explanatory Logic in WBA help to determine whether those causes that are identified by an informal analysis provide a *sufficient* explanation for an incident or accident. Ladkin and Loer introduce the notion of a causal sufficiency criterion [470]. This is based on the argument that for causal relations $A_1 \Rightarrow\!\!\!\succ B, A_2 \Rightarrow\!\!\!\succ B, ..., A_n \Rightarrow\!\!\!\succ B$ the $A_1..A_n$ form a sufficient set of causal factors for B if it would be impossible for B not to happen if $A_1..A_n$ had happened. More formally $A_1..A_n$ form a sufficient set of causal factors for B if and only if:

$$\begin{array}{l} \wedge A_1 \Rightarrow\!\!\!\succ B \\ \wedge A_3 \Rightarrow\!\!\!\succ B \\ \wedge ... \\ \wedge A_n \Rightarrow\!\!\!\succ B \\ \wedge \neg\, B \;\Box\!\!\rightarrow\; \neg\,(A_1 \wedge A_2 \wedge ... \wedge A_n) \end{array} \tag{11.4}$$

From this, Loer goes on to introduce the $\Box\!\!\Rightarrow$ operator to denote both a necessary and sufficient causal relationship. He argues that the goal of the causal sufficiency criterion is to show that:

$$A_1 \wedge A_2 \wedge ...A_n \;\Box\!\!\Rightarrow\; B \tag{11.5}$$

In order to establish such a relationship, analysts can exploit the following rules:

$$\begin{array}{c} C \\ \neg\, C \Rightarrow\!\!\!\succ \neg\, B \\ \neg\, B \Rightarrow\!\!\!\succ \neg\, C \\ \hline C \;\Box\!\!\Rightarrow\; B \end{array} \tag{11.6}$$

$$\begin{array}{l} A \Rightarrow\!\!\!\succ C \\ B \Rightarrow\!\!\!\succ C \\ \hline (A \vee B) \;\Box\!\!\Rightarrow\; C \end{array} \tag{11.7}$$

These rules provide a framework for reasoning about the sufficiency of the semi-formal Why-Because graphs. For example, the left most factors in Figure 11.16 describe a causal relationship between a number of maintenance failures and the initial release of fuel. This relationship can be represents as follows:

$$\begin{array}{l} \mathit{Step}\ 1\ \ (\mathit{Theorem}): \\ \langle \mathit{Forward\ filter\ cover/bolt\ sealing\ surface\ is\ modified\ by\ maintenance} \rangle \wedge \\ \quad \langle \mathit{Copper\ gasket\ is\ deformed\ by\ pressure\ under\ use} \rangle \wedge \\ \qquad [\mathit{Watch\ Engineer\ reuses\ annealed\ copper\ gasket}] \;\Box\!\!\Rightarrow \\ \{\mathit{Watch\ engineer\ fails\ to\ obtain\ fuel - tight\ join\ at\ the\ copper\ gasket} \\ \quad \mathit{sealing\ the\ cover\ to\ its\ securing\ bolt\ on\ the\ forward\ fuel\ oil\ filter}\} \end{array}$$

We can prove this relationship using Loer's meta-rule for deriving the causal sufficiency criterion [499]. The following paragraphs retain the labels that were introduced in the Why-Because graph. This is intended to make the steps of the proof more accessible. Later sections will, however, explain why these annotations might be replaced by predicates with a more precise interpretation. For now it is sufficient to observe that the use of these 'informal' labels makes it difficult to typeset the steps of the proof in a conventional format:

$$\begin{array}{l}
\textit{Step } 2.1 \ \ (\textit{Using } 11.6): \\
\langle \textit{Forward filter cover/bolt sealing surface is modified by maintenance} \rangle \wedge \\
\quad \langle \textit{Copper gasket is deformed by pressure under use} \rangle \wedge \\
\quad\quad [\textit{Watch Engineer reuses annealed copper gasket}]
\end{array}$$

Proof: We can assume that this conjunction is true providing that adequate evidence can be obtained in the aftermath of the incident. The Transportation Safety Board of Canada provided photographic evidence to support these assumptions [621].

$$\begin{array}{l}
\textit{Step } 2.2 \ \ (\textit{Second obligation from } 11.6): \\
\neg\, (\langle \textit{Forward filter cover/bolt sealing surface is modified by maintenance} \rangle \wedge \\
\quad \langle \textit{Copper gasket is deformed by pressure under use} \rangle \wedge \\
\quad\quad [\textit{Watch Engineer reuses annealed copper gasket}]) \mathrel{\Box\!\!\rightarrow} \\
\neg\, \{\textit{Watch engineer fails to obtain fuel} - \textit{tight join at the copper gasket} \\
\quad \textit{sealing the cover to its securing bolt on the forward fuel oil filter}\}
\end{array}$$

$$\begin{array}{l}
\textit{Step } 3.1 \ \ (\textit{Using De Morgan's Law}): \\
\neg\, \langle \textit{Forward filter cover/bolt sealing surface is modified by maintenance} \rangle \vee \\
\neg\, (\langle \textit{Copper gasket is deformed by pressure under use} \rangle \wedge \\
\quad [\textit{Watch Engineer reuses annealed copper gasket}]) \mathrel{\Box\!\!\rightarrow} \\
\neg\, \{\textit{Watch engineer fails to obtain fuel} - \textit{tight join at the copper gasket} \\
\quad \textit{sealing the cover to its securing bolt on the forward fuel oil filter}\}
\end{array}$$

$$\begin{array}{l}
\textit{Step } 4.1 \ \ (\textit{Using } 11.7): \\
\neg\, \langle \textit{Forward filter cover/bolt sealing surface is modified by maintenance} \rangle \mathrel{\Box\!\!\rightarrow} \\
\quad \neg\, \{\textit{Watch engineer fails to obtain fuel} - \textit{tight join at the copper gasket} \\
\quad\quad \textit{sealing the cover to its securing bolt on the forward fuel oil filter}\}
\end{array}$$

Proof: This is true if and only if the engineer obtains a seal in all nearest possible worlds to those in which the forward filter sealing surface is not modified. Given that there was a supply of new gaskets, in other words assuming that the second part of the disjunction in Step 3.1 is false, then the only other way in which the seal could be compromised was through modifications that were not authorised by the manufacturer [621].

$$\begin{array}{l}
\textit{Step } 4.2 \ \ (\textit{Using } 11.7): \\
\neg\, (\langle \textit{Copper gasket is deformed by pressure under use} \rangle \wedge \\
\quad [\textit{Watch Engineer reuses annealed copper gasket}]) \mathrel{\Box\!\!\rightarrow} \\
\neg\, \{\textit{Watch engineer fails to obtain fuel} - \textit{tight join at the copper gasket} \\
\quad \textit{sealing the cover to its securing bolt on the forward fuel oil filter}\}
\end{array}$$

$$\begin{array}{l}
\textit{Step } 5.1 \ \ (\textit{Using De Morgan's Law}): \\
\neg\, \langle \textit{Copper gasket is deformed by pressure under use} \rangle \vee
\end{array}$$

$$\begin{aligned}&\neg\,[\mathit{Watch\ Engineer\ reuses\ annealed\ copper\ gasket}]\ \Box\!\!\rightarrow\\ &\neg\,\{\mathit{Watch\ engineer\ fails\ to\ obtain\ fuel-tight\ join\ at\ the\ copper\ gasket}\\ &\quad \mathit{sealing\ the\ cover\ to\ its\ securing\ bolt\ on\ the\ forward\ fuel\ oil\ filter}\}\end{aligned}$$

$$\begin{aligned}&\mathit{Step}\ 6.1\ \ (\mathit{Using}\ 11.7):\\ &\neg\,\langle\mathit{Copper\ gasket\ is\ deformed\ by\ pressure\ under\ use}\rangle\ \Box\!\!\rightarrow\\ &\neg\,\{\mathit{Watch\ engineer\ fails\ to\ obtain\ fuel-tight\ join\ at\ the\ copper\ gasket}\\ &\quad \mathit{sealing\ the\ cover\ to\ its\ securing\ bolt\ on\ the\ forward\ fuel\ oil\ filter}\}\end{aligned}$$

Proof: This is true if and only if the engineer obtains a seal in all nearest possible worlds to those in which the copper gasket is not deformed under pressure. Additional expert validation is needed to support this argument.

$$\begin{aligned}&\mathit{Step}\ 6.2\ \ (\mathit{Using}\ 11.7):\\ &\neg\,[\mathit{Watch\ Engineer\ reuses\ annealed\ copper\ gasket}]\ \Box\!\!\rightarrow\\ &\neg\,\{\mathit{Watch\ engineer\ fails\ to\ obtain\ fuel-tight\ join\ at\ the\ copper\ gasket}\\ &\quad \mathit{sealing\ the\ cover\ to\ its\ securing\ bolt\ on\ the\ forward\ fuel\ oil\ filter}\}\end{aligned}$$

Conjecture: This is true if and only if the engineer obtains a seal in all nearest possible worlds to those in which the engineer does not reuse an annealed gasket. This theorem is refuted in the following paragraphs.

$$\begin{aligned}&6.3\ \mathit{Q.E.D.\ From}\ 11.7\ \mathit{to}\ 6.1\ \mathit{and}\ 6.2\\ &5.2\ \mathit{Q.E.D.\ From\ De\ Morgan's\ law\ applied\ to\ antecedent\ of}\ 5.1\\ &4.3\ \mathit{Q.E.D.\ From}\ 11.7\ \mathit{to}\ 4.1\ \mathit{and}\ 4.2\\ &3.2\ \mathit{Q.E.D.\ From\ De\ Morgan's\ law\ applied\ to\ antecedent\ of}\ 3.1\end{aligned}$$

$$\begin{aligned}&\mathit{Step}\ 2.3\ \ (\mathit{Third\ obligation\ from}\ 11.6):\\ &\neg\,\{\mathit{Watch\ engineer\ fails\ to\ obtain\ fuel-tight\ join\ at\ the\ copper\ gasket}\\ &\quad \mathit{sealing\ the\ cover\ to\ its\ securing\ bolt\ on\ the\ forward\ fuel\ oil\ filter}\}\ \Box\!\!\rightarrow\\ &\neg\,(\langle\mathit{Forward\ filter\ cover/bolt\ sealing\ surface\ is\ modified\ by\ maintenance}\rangle\wedge\\ &\quad \langle\mathit{Copper\ gasket\ is\ deformed\ by\ pressure\ under\ use}\rangle\wedge\\ &\quad\ [\mathit{Watch\ Engineer\ reuses\ annealed\ copper\ gasket}])\end{aligned}$$

Proof: This is true if and only if the engineer obtains a seal in all nearest possible worlds to those in which the filter is not modified by maintenance and the copper gasket is not deformed under pressure and the engineer does not re-use an annealed copper gasket. If the filter cover had not been modified and the gasket had not been deformed by pressure and been reused then there is no evidence to suggest that the seal would have failed. As before, this argument must be carefully validated by domain experts [195].

$$2.4\ \mathit{Q.E.D.\ From}\ 11.6\ \mathit{applied\ to}\ 2.1,\ 2.2\ \mathit{and}\ 2.3\qquad\Box$$

This proof illustrates how mathematically-based, specification techniques can be used to support the semi-formal structures in a Why-Because graph. As can be seen, the first stage in the proof was to derive a formal representation for the causal relationships that are represented in the left-hand nodes of Figure 11.16. This formalisation provided the theorem that we sought to establish through the use of Loer's meta-rules for the proof of a sufficient causal explanation. The key point here is that these meta-rules provide a template to guide further proofs of the remaining causal relationships in this diagram. Step 1 could be redrafted to formalise these relationships. Steps 2-6 can then be updated. Investigators simply provide the supporting arguments shown for steps 2.1, 4.1, 6.1, 6.2 and 2.3 [499].

This guidance is important because it can help investigators to identify potential weaknesses in their informal reasoning. For example, step 6.2 denoted a causal relationship that is true if and only if the engineer obtains a seal in all nearest possible worlds to those in which the engineer does not reuse an annealed gasket. On closer inspection, it is difficult to defend this argument. Modifications to the seating surface might have compromised the ability of the engineer to achieve a seal even if they had access to a supply of new copper gaskets. Even though we can question this proof step, the the overall proof need not fail. Step 3.1 shows how the argument depends on a disjunction. Step 4.1 has already established the first case and so we need not establish the remainder of the disjunction in order to demonstrate the remainder of the proof. This formal analysis yields several insights. In particular, it illustrates that the annealing of the gasket may not be a necessary cause of the leak. In contrast, the deformation of the gasket and the modifications to the sealing surface together provide the necessary and sufficient causes of the leak, denoted by $\Box\!\!\Rightarrow$.

The previous analysis identified a potential weakness in the previous arguments that have been presented throughout this chapter. The re-use and annealing of the copper gasket need not have been a causal factor in the leak. This argument could prompt investigators to pursue a number of different courses of action. Firstly, they might accept these criticisms and amend the Why-Because graph by omitting the node labelled [Watchkeeping engineer re-uses annealed copper gasket]. Alternatively, further validation might be needed before the results of this formal analysis can be accepted as part of the investigation. This is an important point because incident investigators who are skilled in a particular application domain are unlikely to be familiar with the reasoning techniques that were illustrated in previous pages. In consequence, expert validation is required to support the informal arguments that are made to support the 'Proof' stages for steps 2.1, 4.1, 6.1, 6.2 and 2.3.

The informal arguments that support the previous formal proof are important for a number of reasons. They help to ensure that non-formalists can validate the underlying assumptions that support the formal template or structure that supports the overall causal argument. They also indicate the depth to which investigators want to pursue the formal analysis. The key point here is that it is possible to pursue the formal reasoning beyond the level that was demonstrated in the previous example. For instance, the node (Chief Engineer and Mechanical Assistant do not monitor the port side of the engine room) could be represented by the following clause:

$$\begin{aligned}\neg\ (&attend(chief\_engineer, port\_engine\_room) \vee \\ &attend(mechanical\_assistant, port\_engine\_room))\end{aligned} \tag{11.8}$$

These clauses might then support the extension of formal reasoning techniques from the overall argument structure, shown as the meta-rule given above [499], into the informal arguments that are denoted by the 'Proof' stages for steps 2.1, 4.1, 6.1, 6.2 and 2.3. Ladkin and Loer note that this 'level' of formalisation depends:

> "... on how one wants to analyse the situation; how much one wants to say, what depth and detail of analysis one wants to pursue, the limitations of the language chosen to express the nodes. All of this is very much the choice of the investigator... A similar situation exists in pure axiomatic mathematics. One is provided with sufficient proof rules to get the job done, but what proofs are constructed and how are up to the individual wishes and skill of the user. Proofs may be more detailed or less detailed, easy to follow or cleverly slick, pro forma or creative. Yet the criteria for a valid proof remain constant throughout the enterprise. So with WBA. We have no wish to regulate whether an analysis is most subtle, or how it indicates what future steps to take to prevent recurrences, or whether it must use the latest theory of human-computer interaction. We wish to lay out criteria and reasoning rules for providing a formally-complete causal explanation, according to assumptions that an analyst makes in a particular case. We, thereby make the assumptions clear, explicit and precise, exhibit their role in the explanation, and make the reasoning clear." [470]

The source nodes in a Why-Because graph represent the reasons for an incident or accident. They represent necessary causal factors for an incident or accident. They can easily be identified because they do not have any incoming causal links. This has one very important consequence. Source nodes can be thought of as contingencies that might have be avoided precisely because they lack any necessary causal factors. Table 11.11 summarises the source nodes in Figure 11.16. The rows of this table describe the 'failures' and 'errors' that directly contributed to the incident. They also describe events that might have been appropriate in other contexts. For example, Watchkeeping engineer starts generator need not have caused any problems if everything else had been functioning correctly. However, this event in combination with the failure to obtain a fuel-tight seal led to the initial fire. The compilation of these tables can be used as a further validation for the analytical technique. For example, investigators may be required to justify any decision not to decompose processes into their component factors. The proponents of WBA also argue that source node lists can be used to develop procedures that might avoid particular combinations of adverse events. For example, engineers might be prevented from re-using annealed copper gaskets. Alternatively, maintenance modifications that jeopardise a fuel-tight seal might be closely monitored by supervisory staff.

Factors	Label
State	Forward filter cover and bolt sealing surface is modified by maintenance.
State	Copper gasket is deformed by pressure under use.
Event	Watch engineer re-uses annealed copper gasket.
Event	Watchkeeping engineer starts generator.
Process	Watchkeeping engineer finds all generators and filters are normal.
Process	1st fire party decides to withdraw.
Event	2nd fire party ordered to withdraw.

Table 11.11: Source Node Analysis of Nanticoke WBA Graph

As with the previous analytical techniques in this chapter, it is possible to identify a number of strengths and weaknesses that characterise WBA. For example, the entries in Table 11.11 can be compared to the Transportation Safety Board of Canada's findings about the cause of the Nanticoke incident:

> "The fire was caused by a leakage of fuel, which contacted an exposed exhaust manifold, from the forward fuel filter on the port generator. Contributing to the occurrence was the modification to the fuel filter cover, the re-use of the copper sealing gasket on the cover, the unshielded hot exhaust surfaces adjacent to the filter, and the less-than-adequate engine-room watchkeeping duty during the fire drill before the occurrence." [621]

As can be seen, there is a strong agreement between the informally derived observations of the investigation team and our application of the Why-Because technique. There are, however, a number of important differences. For example, the investigators stressed the significance of the proximity of an exposed ignition source which does not appear as a source node in Figure 11.16. This is a significant omission on our part. The ignition of the fire was represented on the graph as an internal node. We should have added a source state to denote the fact that the indicator tap and the exhaust manifold were exposed. This could have been avoided if the analysts had acquired greater expertise in WBA. It might also have detected during peer review or through a more sustained formal analysis of the causal model. Such omissions are, however, a powerful reminder that even sophisticated analytical techniques are ultimately dependent on the skill and expertise of the individuals who constrict and manipulate the abstractions that they provide.

Having acknowledged the strengths of a traditional 'informal' approach, it is also important to identify potential insights yielded by the more formal style of analysis. Table 11.11 does not simply focus on the causes of the incident itself. It also contains information about the failure of mitigating factors, such as the fire fighting teams. The discipline of listing source nodes can help to check whether the causes of these 'subsidiary' failures are considered in sufficient detail. Table 11.11 helps to reveal, for example, that we have not explained the process by which the first fire party decided to withdraw or the events that led to the order for the 2nd fire party to abandon their work. Additional analysis must be conducted to determine the precise reasons why these attempts were beaten back and, more importantly, whether they were an appropriate response given the state of the fire as it was observed by the crew. This aspect of the incident is, arguably, not considered in sufficient detail by the official report into the incident.

Strauch raises a number of caveats about the application of WBA to the Cali accident [166]. He argued that particular events on a Why-Because graph ought to be distinguished as being more important that others. For example, some decisions have a greater impact on the course of an incident than others. WBA would identify both as 'equal' causes:

> "...not decisions are equal at the time they are made ... each decision alters the subsequent environment, but that while most alterations are relatively benign, some are not. In this accident, this particular decision altered the environment to what became the accident scenario." [763]

These are interesting comments from an individual who has considerable first-hand experience of incident and accident investigations. They could, however, be applied to all of the causal analysis techniques that we have reviewed in this book. The possible exception to this criticism would be the analytical techniques devised to support the application of MORT. As we have seen, investigators can sum the frequency of *what* factors that are associated with *why* nodes to get a raw measure of their relative importance. Weights can also be used to discriminate between the importance of these different failures with common causes. Such techniques suffer from the difficulty of validating any weighting mechanisms that might be used. For instance, how would an investigate discriminate between the relative importance of the deformation of the gasket and the lack of monitoring during the early stages of the fire? Such distinctions are likely to introduce a degree of subjectivity that is intentionally avoided by other aspects of WBA.

There are also a number of deeper philosophical objections to Lewis' use of counterfactual reasoning as it is embodied within WBA. These objections have recently been summarised by Hausman's study of causal asymmetries [313]. Hausman's objections are beyond the scope of this book. Many of his caveats focus on the argument that causes are not counterfactually dependent on their effects. The exposed indicator tap was not counterfactually dependent on the ignition of the Nanticoke fire because the ignition might have been caused by an uncovered exhaust manifold cover. There are possible worlds in which no fire occurred because the exhaust manifold was covered that are at least as similar to the actual world as situations in which a fire did not occur because the indicator was guarded. As we have seen, these situations complicate the application of counterfactual reasoning. Hausman notes that we cannot assume a particular cause simply be observing a set of effects. Each set of effects may be produced by several different causes, even though investigators can identify a determined set of effects for each cause [507]. These observations explain Hausman's choice of 'causal-asymmetries' as the title for his work.

Further criticisms of Lewis' approach focus on the notion of multiple connections. Hausman argues that these occur if a cause d of a is, or in the absence of a, would be connected to b by a path that does not go through a. If there is a multiple connection between a and b, then b will not counterfactually depend on a. Such situations again provide an example of causation without a chain of counterfactual dependence. For instance, the Chief Engineer sounded the general alarm that led the first team of fire fighters to enter the engine room. Their exit caused a second team to be deployed. If we imagine a situation in which the alarm could have led the second team to be deployed whether or not the first had been beaten back then even if we could ensure the success of the first team then there is no guarantee that the second team would not have been deployed. In other words we cannot rely on the argument that if the first team had not been pulled out then the second team would not have been deployed. Both of these caveats affect the other analysis techniques this chapter and Chapter 10 that exploit counterfactual reasoning. It can be argued that these are minor caveats compared to the analytical benefits provided by Lewis' form of reasoning even if, as Hausman argues, 'one cannot defend a counterfactual theory of causation' [313].

The problems of demonstrating the cost-effectiveness of WBA is arguably more important than the theoretical objections proposed by the Hausman's philosophical critique. Semi-formal diagrams, such as Figure 11.16, are relatively cheap and easy to develop. There are some notable differences between this approach and the diagrams employed by MES and STEP. In particular, the ontology of Why-Because graphs including events, states, processes and non-events can be contrasted with the events and conditions of ECF charts. There are, however, considerable similarities. The spatial arrangement of causal relations and the process of informal analysis, including counterfactual reasoning, are comparable. Deeper differences stem from the role of formal reasoning to support the application of Why-Because graphs. These proofs are costly to develop both in terms of the time required and the level of expertise that is essential to guide this process. These formal proofs are important if investigators are to benefit from the strengths of the Why-Because approach. Ladkin and Loer introduce meta-templates that can be used to guide and simplify the formal validation of any causal analysis. Even so, WBA is a time-consuming process. Loer describes a case study during which the development of an 'intuitive' Why-Because graph with approximately 100 nodes required 300 hours. The associated formal proof required a further 1,200 hours [499]. These costs must be assessed against the potential benefits from identifying potential weaknesses in an accident or incident report:

> "We have already been able to identify reasoning mistakes in accident reports using this method. The three accident reports analysed all contained facts which were significantly causally related to the accident, which appear in the WB-graph analysis as causes, but which are not contained in the list of 'probable cause/contributing factors' of the report. We regard this as a logical oversight. (Formally, they appear in the WB-analysis as causal factors that are not themselves explained by any further causal factors; i.e., as source nodes with out-edges but no in-edges.) Some might speculate that there are administrative, political or other social grounds for excluding them from the list of original causal factors, but this is not our interest here. We regard the WB-graph analysis as demonstrating that logical mistakes were made, thereby justifying the use of the WB-analysis to put accident reporting on a rigorous foundation. " [289]

Ultimately, WBA provides many benefits in terms of the precision and rigour that it introduces to causal analysis. Unfortunately, the price that must be paid in order to obtain those benefits is likely to preclude the use of this technique in all but a handful of safety-critical incidents.

This chapter has exploited a deterministic view of the past. We have endeavoured to model a single chain of causal relations that together can help to explain the course of an incident. In our case studies, we have encountered situations where it has not been possible to determine which of a number of possible causal sequences actually led to a mishap. For example, it has not been possible to identify the ignition source in the Nanticoke incident. In general, however, we have attempted to avoid such ambiguity through further investigation. In contrast, the following sections examine ways in which probabilistic models of causation might be applied to support incident and accident analysis. These techniques stem from a scientific and philosophical tradition that questions the notion of deterministic cause [29]. Most of this work has focussed on the problems of using theories of causation as predictive tools. There are, however, important implications for the post hoc use of causal analysis to understand the events the lead to near miss incidents. For example, probabilistic views of causation affect our interpretation of the probability that an accident *might have* occurred. It should be emphasised that the following pages are more speculative than previous sections. We are unaware of any previous attempts to apply these techniques to support incident analysis.

11.3.2 Partition Models for Probabilistic Causation

The previous chapters in this book have assumed that 'causation' can be defined in terms of the necessary and sufficient conditions that must exist between objects in order to achieve particular effects. In particular, counterfactual arguments have been used to identify situations in which a set of effects would not have occurred if those necessary and sufficient conditions had not been fulfilled. It is important to note that a number of caveats can be raised to these general theories of causation.

For instance, previous sections have identified different forms of causal asymmetry. For instance, if necessary and sufficient conditions do not hold then an effect may still occur. This complicates the application of counterfactual argumentation when investigators use a form of 'backtracking' to identify causes from their effects. Similarly, many physicists maintain that occurrences are not determined [200]. In other words, we can never be absolutely certain that a set of effects will be produced even if necessary and sufficient conditions can be demonstrated to hold at a particular moment. In contrast, it is argued that a complete specification of the state of a system only determines a set of probabilities [313]. Some of the proponents of this view have argued that what happens in any given situation owes as much to chance as it does to cause. This analysis has profound implications. For instance, we might be persuaded to abandon the notion of 'sufficient' causes that do not account for this role of chance! In this view, causal analysis would owe more to probabilistic risk assessment and human reliability assessment than it does to the discrete mathematics of WBA or Causal Trees. This is an interesting conjecture. Such an approach might emphasise the role of performance shaping factors incident rather than discrete events [443]. Instead of focusing on the identification of a deterministic sequence of cause and effect relationships, which are difficult to validate given the problems of causal asymmetry mentioned above, investigators should focus on those conditions that made effects more likely within a given context. For instance, we might describe the Nanticoke incident in terms of the probabilities that either the indicator tap or the manifold ignited the fire.

It is important to emphasise, however, that probabilistic forms of analysis do not eliminate the need to consider causality. For example, supposing that a factory produced a faulty gasket and that this gasket eventually led to a fuel leak on board a ship. Investigators might argue that the gasket caused the leak even though the production of the gasket created a small probability that any particular vessel would be affected. Statistical mechanics has also identified mass populations for which particular relations are deterministic, however, the best means of describing mass effects is through the use of probabilistic techniques [313]. This is important within the field of incident analysis because, as we shall see, national reporting systems typify these mass phenomena. For instance, we might receive ninety-nine reports in which a fire is caused by the exposure of an ignition source to a fuel supply. In one report, however, the same circumstances might not have led to a fire. Although we have an apparently deterministic model of how a fire starts, there may be exceptions that persuade analysts to consider probabilistic aspects of causation. These exceptions characterise many different aspects of incident analysis and, more generally, of individual attitudes to causation. For instance, people often argue that fines cause reductions in health and safety violations even though they do not believe that the deterrence is perfect. Similarly, people will say that dropping a glass causes it to break even though they have seen similar situations in which the glass did not break. It is often argued that a more complete knowledge of the moment acting on the glass would enable causal explanations of why certain glasses break while others do not. However this indeterminism is equally apparent in the 'microscopic' causal relations that explain the physics of these different outcomes.

Probabilistics approaches to causal analysis raise many practical and theoretical questions. The frequentist approach derives the probability of an event from an analysis of comparative frequencies. We can use information about previous fires to derive numerical estimates for the number of times that ignition was caused by a manifold or by an exposed indicator tap. Previous sections have dismissed this approach because it can be difficult to validate the frequency of rare events. We shall return to this theme several times in the following pages. Alternatively, empirical analysis can be used to repeatedly recreate situations in which either of these sources might ignite leaking oil. Again, frequencies can be calculated to derive probability estimates. This approach raises questions about the validity of the experimental context in which the simulations are conducted.

Unfortunately, a number of factors complicate the use of probabilistic approaches to causal analysis. Raw event frequencies cannot, typically, be used to determine the probability of particular 'causes' in the aftermath of an incident. For example, an examination of previous fires might find that six were caused by indicator taps and ten were caused by exposed manifolds. Supposing, however, that nine of the ten manifold fires involved a different fuel leak than that on the Nanticoke. In this situation, any causal analysis must draw upon conditional probabilities. These represent the probability of an event given that some other factors hold. In this case we need to know the probability of ignition from each source given the fuel leak characteristics that held during the Nanticoke fire. This use of conditional probabilities has some significant benefits for incident analysis. Investigators are not dealing with prior probabilities describing future events where we know relatively little additional information about the potential state of a system. In the aftermath of an incident it is often possible to obtain the conditioning information that helps to support particular probability assessments. The following section, therefore, extends this analysis to consider Bayesian statistics. For now it is sufficient to observe that these techniques can be used to represent and reason about a hypotheses given particular evidence in the aftermath of an incident.

As mentioned, an important limitation of many probabilistic approaches to causation is that it can be difficult to validate numerical estimates of rare events. Fortunately, many probabilistic theories of causation avoid this problem by describing how particular causes make their effects 'more likely'. For instance, Hempel argues that a and b are causally connected in a context C if there is a very high probability that b is true given that a is true in C: $Pr(b \mid a \wedge C)$ [345]. For instance, we could say that the maintenance modification to the sealing surface of the Nanticoke's fuel filter, a, was a cause of the leak, b, because this modification made the leak very likely given everything else that was discovered about the incident including the failure to report such problems etc, C. Hempel's approach also avoids the need to assign precise numeric values to individual probabilities it also creates the problem that investigators must determine what is meant by 'very likely'. It is possible, however, that this theoretical objection can be addressed by experience in applying the technique within a particular domain. The following paragraphs explain how Hempel's ideas might contribute to a method for the causal analysis of adverse incidents:

1. *Record the context in which an incident occurs.* This step ensures that as much information as possible is derived from the primary and secondary investigation of an incident. Previous sections have mentioned the difficulty of predicting all of the information that might be relevant to a causal analysis and so investigators should collate as much data as possible. Chapter 15 will examine the practical problems that such a policy creates for information storage.

2. *Perform an initial deterministic causal analysis.* Having collated as much information about the context, C, in which an incident occurs, investigators can exploit one of the causal analysis techniques introduced in previous sections. For instance, STEP or WBA might be used to identify potential causal factors in the immediate aftermath of an incident. Chapter 12 will describe how these techniques can be used to derive initial recommendations that are intended to avoid any recurrence of an near-miss occurrence.

3. *Build up sufficient data to perform a statistical analysis of potential causes.* Over time an incident reporting system may gather information about a number of adverse occurrences that have similar outcomes, b. Investigators can then examine the contextual information that has been recorded for each incident, C, to identify those events, a, that have the highest relative frequency. These events need not, however, have any causal relationship to b. For instance, b might occur before a in the temporal ordering of events. Additional techniques, such as WBA, must therefore validate the causal relations that are induced by the statistical analysis of incident collections. This form of causal analysis does, however, avoid the bias that can arise from causal asymmetries. Analysts do not simply use deterministic models to search for a narrow range of causes that can be made to 'fit' the observed effects.

The approach, described above, has numerous potential benefits from its integration of deterministic and probabilistic models of causation. The initial use of deterministic approachs can help to direct resources to a number of clearly defined causal factors in the aftermath of an incident. Probabilistic techniques can be used to search for other causal factors through an analysis of the correlations that exist between common factors in similar incidents. As far as we are aware, this approach has not been explicitly described before. It is, however, increasingly being adopted by many commercial and regulatory organisations. Chapter 15 will describe how probabilistic information retrieval tools have been developed to exploit correlations between the terms that are used to describe both the consequences and the causes of incidents and accidents.

As mentioned, Hempel's initial formulation provided little guidance on the meaning of the term 'very likely'. Fortunately, a number of refinements have been made to these early ideas. One of these approaches holds that a is causally related to b in a context C if the probability of A and B in C is not the same as the probability of B in C and the probability of A in C:

$$Pr(B \wedge A \mid C) \neq Pr(B \mid C).Pr(A \mid C) \tag{11.9}$$

We assume that we cannot derive A or $\neg A$ from C. Upper case denotes types, lower case is used to denote tokens of a particular type; token a is of type A and so on. This inequality has some interesting properties that can be applied to guide the causal analysis of incidents and accidents. Recall from Chapter 9 that $Pr(a \wedge b) = Pr(a).Pr(b)$ depends upon the independence of both a and b. If there is a causal connection between A and B then we might expect that the occurrence of a would make b more likely. Conversely, if A is a barrier to B then an occurrence of a will make b less likely. Hausman argues that a is positively causally related to b when the probability of A and B given C is greater than the probability of B given C multiplied by the probability of A given C [313]. In other words, a causal relationship implies that the probability of there being a general fire alarm, a, and a Halon system being deployed, b, on board a vessel, C, is greater that the probability of a general fire alarm being issued multiplied by the probability of a Halon system being deployed in similar circumstances. The deployment of the Halon system might be a relatively rare event compared to the sounding of a general alarm. However, a causal relationship with the alarm might result in a much higher probability being associated with situations in which the alarm and the Halon deployment both occur than situations in which we only know that one of these events has occurred:

$$Pr(B \wedge A \mid C) > Pr(B \mid C).Pr(A \mid C) \tag{11.10}$$

The key point to understanding this formula is that causes do not make their effects probable. They simply make them more probable than they otherwise would have been. We can also say that a is negatively causally related to b when the probability of A and B given C is less than the probability of B given C multiplied by the probability of A given C [313]. For instance, the probability of an engineer failing to obtain a fuel-tight seal, b, and of that engineer reporting the problem associated with the sealing surface, a, are together less than the independent probabilities of the engineer reporting the problem multiplied by the probability of the engineer failing to obtain the seal. This follows because the fact that the engineer reported the maintenance problem makes it less likely that they will be satisfied by any subsequent attempt to form a seal on the damaged surface:

$$Pr(B \wedge A \mid C) < Pr(B \mid C).Pr(A \mid C) \tag{11.11}$$

From this line of argument, we can say that a's cause b's under circumstances C if a's precedes b's in the temporal sequence leading to an incident and it is the case that the probability of B and A in C is greater than the probability of B given that we know $\neg A$ and C. Or we can say that a's cause b's under C if a's precedes b's in the temporal sequence leading to an incident and it is the case that the probability of B and A in C is greater than the probability of B given only C:

$$\begin{gathered} Pr(B \wedge A \mid C) > Pr(B \mid \neg A \wedge C) \vee \\ Pr(B \mid A \wedge C) > Pr(B \mid C) \end{gathered} \tag{11.12}$$

Unfortunately, this formalisation leads to further problems. For example, it may be that a precedes b and that $Pr(B \wedge A \mid C) > Pr(B \mid \neg A \wedge C)$ but that a and b and effects of the *same* cause. One way to avoid this is to examine the events prior to a to determine whether there is another event that might 'screen off' or account for both a and b. Further models have been developed to formalise this approach [765] and these, in turn, have been further criticised [313]. The key point here is to provide an impression of the complexity that must be address by any attempt to exploit probabilistic models of causation as a means of supporting incident analysis. The initial appeal of an alternative to deterministic models rapidly fades as one considers the complexity of an alternative formulation.

One important source of additional complexity is that causal factors may both promote and confound particular effects. In this refinement, some factor that causes as to occur can have an independent negative influence on the occurrence of b's. For instance, the probability that the Nanticoke fire would lead to the loss of the vessel was increased by the lack of effective monitoring when the initial fire developed on the port side of the engine room. This lack of monitoring might have been a result of having both the Chief Engineer and the Mechanical Assistant in the Control Room during the fire drill. However, the same circumstances that interfered with their monitoring responsibilities may also have reduced the probability that the fire would jeopardise the safety of the vessel because both crewmembers could initiate the eventual response to the incident. Similarly, the increasing probability of b from a by one causal path can be offset by negative influences from a along another causal path. For example, the fire drill procedures may have made it more likely that the vessel would be seriously damaged by distracting members of the crew from their normal activities. The same drills may have made it less likely that the vessel would be seriously damaged because members of the crew were already prepared to respond to the general alarm that was sounded by the Chief Engineer. The importance of these mitigating factors has been repeatedly emphasised in recent studies of incident reporting systems [842]. Unfortunately, these factors are not adequately represented within many deterministic causal analysis techniques.

The proponents of probabilistic theories of causation have responded to these observations by revising the previous formulations to include a partition S_j of all relevant factors apart from A and C. From this it follows that a's cause b's in circumstances C if and only if:

$$\forall j : Pr(B \mid A \wedge S_j \wedge C) > Pr(B \mid S \wedge C) \tag{11.13}$$

$\{S_j\}$ is a partition of all relevant factors excluding A and C. These factors represent the negative or positive causal factors, $c_1, ..., c_m$, that must be held fixed in order to observe the causal effect of a. We require that any element, d, of a subset in S_j is in c_i if and only if it is a cause of b or $\neg\, b$, other than a, and it is not caused by a. For instance, a hot manifold is liable to have a negligible impact on an existing fire. We can, therefore, include a factor, c_i, in each subset to require that a fire must not have already started in order for a hot manifold, a, to ignite a fuel source, b. Each of the factors in $c_1, ..., c_m$ must be represented in each subset. Each factor must also either be present or absent; there may or may not be an existing fire. This results in 2^m possible combinations of present or absent factors. Some combinations of the factors $c_1, ..., c_m$ will be impossible. Hence some combinations of c_i can be excluded from S_j. For example, it is difficult to foresee a situation in which the engine room is flooded with Halon gas and the fire continues to burn. Yet both of these factors could prevent us from observing an ignition caused by a hot manifold. Other combinations may result in b being assigned a probability of 1 or 0 regardless of a. For instance, if the engine room were flooded with Halon then the fire should not ignite irrespective of the state of the exhaust manifold. As mentioned, these impossible combinations and combinations that determine b are omitted from S_j. All the remaining combinations of causal factors must be explicitly considered as potential test conditions and are elements of S_j. In other words, a's must cause b's in every situation described by S_j.

Some proponents of this partition theory dispense with any explicit representation of the context, C [153]. This approach relies entirely upon the partitioning represented by S_j. This is misleading. Causal relations may change from one context to another. For instance, the effects of a fuel leak may depend upon the pressure at which the fuel escapes. This, in turn, may depend upon the size and configuration of a generator. The meta-level point here is that we would like causal relations to hold over a variety of circumstances, these are characterised by S_j. We cannot, however, expect to identify causal relations that are not relativised to some background context [313].

A number of objections have inspired further elaborations to this partition model of causation [221]. In terms of this book, however, we are less interested in the details of these reformulations than we are in determining whether these models might support the causal analysis of incidents. The abstract model, outlined above, provides a structure for the analysis of incidents in the same way that Why-Because graphs and the associated proof templates provided by Ladkin and Loer also provide a structure for causal analysis. For example, we can apply the partition model to the Nanticoke example by identifying candidate causal relations. Investigators can use their domain expertise to determine those relations that are then subjected to a more formal analysis, this equates to stage 2 of the method proposed for Hempel's model given above. For instance, previous sections have argued that it is difficult to determine the ignition source for the Nanticoke fire. This causes problems for deterministic causal models. We might, therefore, exploit the partition model to represent a causal relationship between an ignition event, b, and the fuel oil coming into contact with an exhaust manifold. As mentioned, C represents all state descriptions for the system under consideration. We might, therefore, informally argue that C represents the state of any merchant vessel that relied upon diesel generators. This context might be narrowed if the formalisation of the incident is intended only to apply to a restricted subset of these ships. In contrast, it might be extended if the formalisation also captures important properties of other vessels, such as military ships that employ diesel generators. Irrespective of the precise interpretation, it is important that analysts explicitly identify this context that helps other investigators to understand the scope of the model. We can then go on to identify other causal factors that might be represented in subsets of the form $c_1, ..., c_m \in S_j$. Recall that d is in $c_1, ..., c_m$ if and only if it is a cause of b or $\neg\, b$, other than a, and it is not caused by a:

c_1 represents 'the room floods with Halon',
c_2 represents 'fuel is sprayed at pressure',
c_3 represents 'shielding protects the manifold'.

As mentioned, individual factors may either be present or absent during particular incidents. There are, therefore, 2^3 potential elements of S_j. In the following, the omission of an element from any set implies that the causal factors are omitted. The first sequence represents a situation in which all of the previous causal factors are present. The room floods with Halon and the fuel is sprayed at pressure and shielding protects the manifold. The second of the subsets indicates that all of the factors are true except for the last one; the shielding does not protect the manifold.

$$\{c_1, c_2, c_3\}, \{c_1, c_2\}, \{c_1, c_3\},$$
$$\{c_2, c_3\}, \{c_1\}, \{c_2\}, \{c_3\}, \{\},$$

We can, however, reduce the number of combinations that we need to consider in order to establish a causal relation between a and b. As mentioned, some combinations of these causal factors are impossible. Other combinations may entirely determine the effect irrespective of the putative cause. For example, we can ignore any subset that contains c_1. If the room floods with Halon then the fire will not ignite, b, whatever happens to the fuel and the manifold, a. Conversely, we can insist that all subsets must include c_2. If the fuel is not sprayed at pressure then the fire will not ignite even if the fuel oil comes into contact with an exhaust manifold; as the manifold may not reach the flash-point of the fuel. In order to establish causality, we must however consider whether a increases the probability of b taking all other combinations of the causal factors, c_i, into account:

$$\{c_2, c_3\}, \{c_2\}.$$

In other words, in order for a causal relation to hold between between an ignition event, b, and the fuel oil coming into contact with an exhaust manifold, a, we must show that the effect would still be more likely if fuel is sprayed at pressure whether or not shielding protected the manifold. The shielding might reduce the absolute probability of the ignition but may not necessarily reduce it to zero, as a Halon deployment might. We must, therefore, show that the cause still increases the probability of the effect in both of these conditions.

This application of the partition model has a number of practical advantages. For instance, investigators are not forced to quantify the probability that a cause will yield a particular effect. The partition model also offers some advantages when compared to more deterministic models. This approach provides an elegant means of dealing with uncertainty about the precise causes of an incident. In particular, previous analyses have experienced acute problems from the investigators difficulty in determining what caused the ignition of the fire on the Nanticoke. The partition model entirely avoids this problem. It is possible to characterise multiple potential causes using the relevant factors represented by c_i. For example, we could have extended C_i to include fuel oil comes into contact with an indicator tap. We can also use the same techniques to represent and reason about the impact of mitigating factors. This again was problematic in deterministic techniques. In the previous example, we had to demonstrate that an ignition was more likely to occur whether or not the manifold was protected by shielding. We also showed how the same approach can represent barriers, such as Halon deployment, that prevent an effect from occurring. It is important to stress that these arguments about the probability of an ignition must be validated [195]. The partition model helps here because analysts can explicitly represent the anticipated impact of contributory causes, of mitigating events and of potential barriers. In contrast, many deterministic techniques consider these issues as secondary to the process of enumerating those failures that led to an incident.

There are, however, a number of practical concerns that arise during the application of the partition model of non-deterministic causation. All of the relevant factors, c_i, in the previous example were carefully chosen to be events. This satisfies the requirement that 'd is in $c_1, ..., c_m$ if and only if it is a cause of b or $\neg$ b, other than a, and it is not caused by a'. Previous informal examples in this section have argued that a hot manifold would not have ignited the fire if a fire had already been burning. Ideally, we would like to extend c_i to include an appropriate state predicate so that we can explicitly represent and reason about such a situation. Alternatively, we could refine the relatively abstract view of the context, C, that was introduced in this example. Further concerns stem from the problems of applying an abstract model of causation to support incident analysis. It is entirely possible that the previous example reflects mistakes in our interpretation of the theoretical work of Cartwright [153] and Hausman [313]. Further work is, therefore, needed to determine appropriate formulations and interpretations of these non-deterministic models. This brief example does, however, demonstrate that probabilistic approaches can avoid some of the problems that uncertainty creates for the deterministic techniques that have been presented in previous sections.

There are also a number of more theoretical concerns about the utility of partition models. The formula (11.13) ensures that a increases the probability of b irrespective of the values assigned to these other relevant factors. This provides a definition of causation in which the mitigating effects of these relevant factors must not offset the increased probability of an associated effect. This might seem a reasonable criterion for causality. It does, however, lead to a number of philosophical problems. For instance, it might be argued that the crew's failure to regularly inspect the engine room is a potential cause of major fires such as that on board the Nanticoke. It might equally be argued that, under certain circumstances, regular inspection of the engine room might lead to a major fire. For example, operators might miss an automated warning in the control room that indicated a potential problem elsewhere in the engine room [621]. Under the system described above, such circumstances would prevent investigators from arguing that lack of inspection is a cause of major fires! This is a general problem; there are contexts in which "smoking lowers one's probability of getting lung cancer, drinking makes driving safer and not wearing seat-belts makes one less likely to suffer injury or death" [313]. As before there are a number of refinements on the basic model outlines in (11.13). It remains to be seen whether any of these extensions might provide an adequate framework for the causal analysis of safety-critical incidents. It might seem to be far-fetched that probabilistic models of causation might yield pragmatic tools for incident analysis. Against this one might argue that Lewis' possible worlds semantics for counterfactual reasoning would have appeared equally arcane before the development of WBA.

11.3.3 Bayesian Approaches to Probabilistic Causation

There are a number of alternative semantic interpretations for the Pr function introduced in the previous section [150]. In particular, Pr may be viewed either as a measure of confirmation or as a measure of frequency. The former interpretation resembles the Bayesian view; probability is contingent upon the observation of certain evidence. The latter resembles the manner in which engineers derive reliability figures. Estimates of pump failures are derived from the maintenance records of plant components. This distinction has been a subject of some controversy. For instance, Carnap argued:

> "... for most, perhaps for practically all, of those authors on probability who do not

> accept a frequency conception the following holds. i. Their theories are objectivist (and) are usually only preliminary remarks not affecting their actual working method. ii. The objective concept which they mean is similar to (the frequency view of) probability." [150]

Brevity prevents a detailed explanation of the contrasting positions in this debate. It is, however, possible to illustrate the common origin for these two different approaches. Both the partition models and Bayesian views exploit conditional probabilities. These also formed the foundation for the treatment of probabilistic causality in the previous chapter. As before, we use the following form to denote that the probability of the event B given the event A in some context C is x.

$$Pr(B \mid A \wedge C) = x \tag{11.14}$$

From this we can derive the following formula, which states that the conditional probability of B given A in C multiplied by the probability of A in C is equivalent to the probability of A and B in C. In other words the probability of both A and B being true in a given context is equivalent to the probability of A being true multiplied by the probability that B is true given A:

$$Pr(B \mid A \wedge C).Pr(A \mid C) = Pr(A \wedge B \mid C) \tag{11.15}$$

We can use this axiom of probability calculus to derive Bayes theorem:

$$Pr(B \mid A \wedge C).Pr(A \mid C) = Pr(B \wedge A \mid C) \qquad (\textit{Commutative Law applied to } \wedge \textit{ in } (11.15)) \tag{11.16}$$

$$Pr(B \mid A \wedge C).Pr(A \mid C) = Pr(A \mid B \wedge C).Pr(B \mid C) \qquad (\textit{Substitution of RHS using } (11.15)) \tag{11.17}$$

$$Pr(B \mid A \wedge C) = \frac{Pr(A \mid B \wedge C).Pr(B \mid C)}{Pr(A \mid C)} \qquad (\textit{Divide by } Pr(A \mid C)) \tag{11.18}$$

The key point about Bayes' theorem is that it helps us to reason about the manner in which our belief in some evidence affects our belief in some hypothesis. In the previous formula, our belief in B is affected by the evidence that we gather for A. It should be emphasised that (11.18) combines three different types of probability. The term $Pr(A \mid C)$ represents the *prior* probability that A is true without any additional evidence. In the above, the term $Pr(B \mid A \wedge C)$ represents a *posterior* probability that B is true having observed A. We can also reformulate (11.18) to determine the *likelihood* of a potential 'cause' [200]. The following formula considers the probability of a given hypotheses, B, in relation to a number of alternative hypotheses, B_i where B and B_i are mutually exclusive and exhaustive:

$$Pr(B \mid A \wedge C) = \frac{Pr(A \mid B \wedge C).Pr(B \mid C)}{Pr(A \mid B \wedge C).Pr(B \mid C) + \sum_i Pr(A \mid B_i \wedge C).Pr(B_i \mid C)} \tag{11.19}$$

The previous formula can be used to assess the likelihood of a cause B given that a potential effect, A, has been observed. This has clear applications in the causal analysis of accidents and incidents. In particular, (11.19) provides a means of using information about previous incidents to guide the causal analysis of future occurrences.

In the Nanticoke case study, investigators might be interested to determine the likelihood that reported damage to an engine room had been caused by the pressurised release of fuel from a filter. The first step would involve an analysis of previous incidents. This might reveal that fuel from a filter was identified as a cause in 2% of previous mishaps. Lubrication oil might account for 1%. Other fuel sources might together account for a further 3% of all incidents:

$$Pr(\mathit{filter\ fire} \mid C) = 0.02 \tag{11.20}$$

$$Pr(\mathit{lube\ fire} \mid C) = 0.01 \tag{11.21}$$

$$Pr(\mathit{other\ fire} \mid C) = 0.03 \tag{11.22}$$

In order to gain more evidence, investigators might try to determine how likely it is that one of these fires would cause serious damage to an engine room. Further analysis might reveal that thirty per cent of previous incidents involving the ignition of filter fuel resulted caused significant damage to an engine room. Twenty per cent of lube fires and fifty per cent of fires involving other fuel sources might have similar consequences:

$$Pr(\mathit{engine\ room\ damage} \mid \mathit{filter\ fire} \wedge C) = 0.3 \tag{11.23}$$

$$Pr(\mathit{engine\ room\ damage} \mid \mathit{lube\ fire} \wedge C) = 0.2 \tag{11.24}$$

$$Pr(\mathit{engine\ room\ damage} \mid \mathit{other\ fire} \wedge C) = 0.5 \tag{11.25}$$

We can now integrate these observations into (11.19) to calculate the probability that a filter fuel fire was a cause given that a serious engine room fire has been reported. This following calculation suggests that there is a twenty-six per cent chance that such a filter fire had this effect:

$$\begin{aligned} & Pr(\mathit{filter\ fire} \mid \mathit{engine\ room\ damage} \wedge C) \\ & = \frac{Pr(\mathit{engine\ room\ damage} \mid \mathit{filter\ fire} \wedge C).Pr(\mathit{filter\ fire} \mid C)}{((Pr(\mathit{engine\ room\ damage} \mid \mathit{filter\ fire} \wedge C).Pr(\mathit{filter\ fire} \mid C))} \\ & \quad + (Pr(\mathit{engine\ room\ damage} \mid \mathit{lube\ fire} \wedge C).Pr(\mathit{lube\ fire} \mid C)) \\ & \quad + (Pr(\mathit{engine\ room\ damage} \mid \mathit{other\ fire} \wedge C).Pr(\mathit{other\ sources} \mid C)) \qquad (11.26) \end{aligned}$$

$$= \frac{(0.3).(0.02)}{(0.3).(0.02) + (0.2).(0.01) + (0.5).(0.03)} \tag{11.27}$$

$$= 0.26 \tag{11.28}$$

A number of caveats can be raised against this application of Bayes' theorem. Many concerns centre on our use of evidence about previous mishaps to guide the causal analysis of new incidents. The previous calculations relied upon investigators correctly identifying when a fire had caused 'significant damage' to an engine room. These is a danger that different investigators will have a different interpretation of such terms. Chapter 15 describes how Bayesian techniques can account for the false positives and negatives that result from these different interpretations. For now it is sufficient to observe that our analysis of previous incident frequencies might bias the causal analysis of future incidents. For instance, we have made the assumption that these incidents occurred in comparable contexts, C. There may be innovative design features, such as new forms of barriers and protection devices, that would invalidate our use of previous frequencies to characterise future failures.

Dembski argues that it is seldom possible to have any confidence in prior probabilites [200]. Such figures can only be trusted in a limited number of application domains. For instance, estimates of the likelihood of an illness within the general population can be validated by extensive epidemiological studies. It is difficult to conduct similar studies into the causes of safety-critical accidents and incidents. In spite of initiatives to share incident data across national boundaries, there are few data sources that validate the assumptions represented in (11.23), (11.24) and (11.25). This book has identified a number of different biases that affect the use of data from incident reporting systems. For example, Chapter 5 referred to the relatively low participation rates that affect many incident reporting schemes. This makes it difficult for us to estimate the true frequency of lube fires or filter fires. These incidents may also be extremely rare occurrences. It can, therefore, be very difficult for investigators to derive the information that is required in order to apply (11.19).

In the absence of data sources to validate prior probabilities, investigators typically rely upon a variant of the indifference principle. This states that given a set of mutually exclusive and exhaustive possibilities, the possibilities are considered to be equi-probable unless there is a reason to think otherwise. This would lead us to assign the same probabilities to fires being caused by filter fuel, to lube oil and to all other sources. Unfortunately, the pragmatic approach suggested by the indifference principle can lead to a number of paradoxes [369]. Objections can have also been raised against any method that enables investigators to move from conditional probabilities, such as $Pr(A \mid B_i \wedge C)$, to their 'inverse' likelihoods, $Pr(B_i \mid A \wedge C)$ [200].

The use of subjective probabilities provides a possible solution to the lack of frequential data that might otherwise support a Bayesian approach to the causal analysis of safety-critical incidents. Subjective probabilities are estimates that individual investigators or groups of investigators might make about the probability of an event. For example, a subjective probability might be an individuals assessment of the chances that lube oil could start a fire that might cause serious damage to an engine room. One standard means of estimating this probability is to ask people to make a choice between two or more lotteries. This technique is usually applied to situations in which it is possible to associate financial rewards with particular outcomes. Von Neumann and Morganstern provide a detailed justification for the general applicability of this approach [626]. Certain changes must, however, be made in order to explain how these lotteries might support the causal analysis of adverse incidents.

1. I might be offered a situation in which there is a certainty that if a lube oil fire occurs in the next twelve months then it will result in major damage to an engine room;

2. alternatively, I might be offered a form of 'gamble'. This requires that I select a token at random from a jar. This jar contains N tokens that are marked to denote that there is no serious damage to an engine room during the next twelve months. The remaining 100-N tokens are marked to denote that there is such an incident.

I will prefer option (2) if every token indicates that engine rooms remain undamaged, i.e. N=100. Conversely, I will prefer option (1) if every token indicates the opposite outcome, i.e. N=0. This requires additional explanation. Recall from option (1) that engine room damage will occur *if* there is a lube oil fire. In option (2), if N=0 then there is a certainty that engine room damage will occur. This explains the preference for (1), the individual makes a subjective assessment of the likelihood of the lube fire and then must trade this off against the potential for there not to be engine room damage in (2). There will, therefore, be a value of N for which the two situations are equally attractive. At such a position of indifference, $\frac{N}{100}$ is my estimate of the probability that a lube oil fire will cause serious damage to an engine room in the next year. Jensen notes that "for subjective probabilities defined through such ball drawing gambles the fundamental rule can also be proved" [398]. This fundamental rule is given by formula (11.15) that provided the foundation for Bayes' theorem (11.18).

A number of problems affect the use of subjective probabilities. An individuals' preference between the two previous options is not simply determined by their subjective estimate of the probability of a lube oil fire. It can also be affected by their attitudes towards risk and uncertainty. For example, one person might view that a 20% chance of avoiding engine room damage is an attractive gamble. They might, therefore, be willing to accept N=20. Another individual might be very unwilling to accept this same gamble and might, therefore, prefer option (1). These differences need say very little about the individual's view of the exact likelihood of a lube fire resulting in major engine room damage. In contrast, it may reveal more about their attitude to uncertainty. The first individual may choose the gamble because they have more information about the likelihood of engine room damage than they do about the lube fire in (1). Individual preferences are also affected by attitude to risk [689]. Experimental evidence has shown that different individuals associate very different levels of utility or value to particular probabilities. A risk adverse individual might view a 20% gamble as a very unattractive option whereas a risk preferring individual might be more inclined to accept the risk given the potential rewards.

In spite of the problems if deriving both frequentist and subjective probabilities, Bayesian inference has been used to reason about the dependability of hardware [86, 294] and software systems

[497]. In particular, a growing number of researchers have begun to apply Bayesian Networks as a means of representing probabilistic notions of causation. it is based around the concepts on contingent probability that, as we have seen, can also arguably be used to provide insights into the likelihood of certain causes. Figure 11.18 presents a Bayesian network model for one aspect of the Nanticoke incident. Investigators initially observed horizontal soot patterns on top of valve covers 4, 5 and 6 and a shadowing on the aft surfaces of these structures. These observations indicate that the fire originated on the port side of the engine, forward of cylinder head number 1. These effects might have been caused by a fire fed from one of two potential sources. This is indicated in Figure 11.18 by the two arrows pointing into the node labelled horizontal soot patterns.... The arrows point from a cause towards the effect. The + symbols indicate the cause makes the effect more likely. Conversely, a barrier might be labelled by a - symbol if it made an effect less likely. As can be seen, the two potential causes of the horizontal soot patterns include a lube oil leak from under that valve cover near cylinder head number 6. These effects might also have been caused by a fuel oil leak from one of the filters. Further investigations reveal that the valve covers were in tact and in place after the fire. This increases the certainty that the fire started from a filter leak rather than a lube oil leak under the valve covers. Another way of looking at Figure 11.18 is to argue that leaks from either the filter or from lube oil are consistent with the horizontal soot patterns. Only a fuel oil leak from the filter is consistent with the valve covers being in tact after the fire.

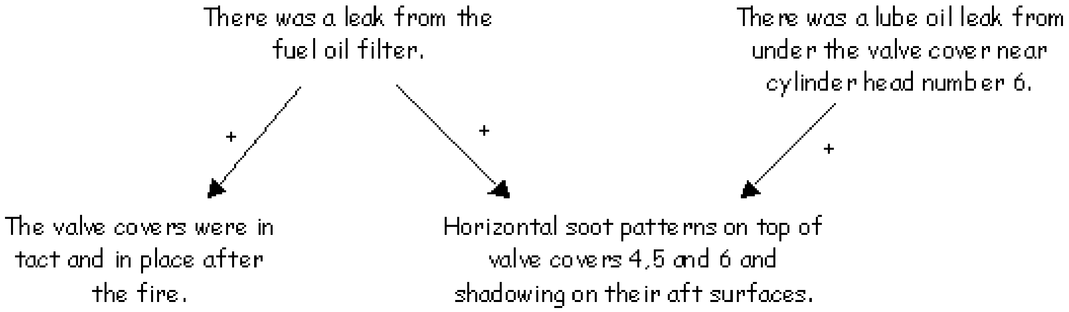
There was a leak from the fuel oil filter.
There was a lube oil leak from under the valve cover near cylinder head number 6.
+
+
+
The valve covers were in tact and in place after the fire.
Horizontal soot patterns on top of valve covers 4,5 and 6 and shadowing on their aft surfaces.

Figure 11.18: Bayesian Network Model for the Nanticoke Fuel Source

Before continuing to apply Bayesian networks to support our causal analysis of the Nanticoke incident, it is important to observe that some authors have argued that these diagrams must not be used as causal models. In contrast, they should only be used to model the manner in which information propagates between events. This caution stems from doubts over methods that enable investigators to move from conditional probabilities to their 'inverse' likelihoods, mentioned in previous paragraphs [200]. This point of view also implies further constraints on the use of Bayesian networks. For instance, it is important not to model interfering actions within a network of information propagation. Jensen provides a more complete introduction to these potential pitfalls [398].

	filter fire	¬ *filter fire*
valve covers ok	1	0.98
¬ *valve covers ok*	0	0.02

Table 11.12: Conditional Probabilities for the Bayesian Analysis of the Nanticoke Incident (1)

The quantitative analysis of Figure 11.18 begins with either a frequentist or subjective estimate of the likelihood of each cause. Recall that $Pr(\textit{filter fire} \mid C) = 0.02$ and that $Pr(\textit{lube fire} \mid C) = 0.01$. We can use these prior probabilities and the information contained in Figure 11.18 to derive the conditional probabilities for $P(\mathit{filter\ fire} \mid \mathit{valve\ covers\ ok})$. These are shown in Figure 11.12. If we know that there was a filter fire then it is certain that the valves would be in tact, this represents a simplifying assumption that can be revised in subsequent analysis. If there was not a filter fire then there is a 0.98 probability that the valves would be in tact but a 0.02 probability that they would not. The conditional probabilities shown in Figure 11.12 are represented in matrix form throughout the remainder of this analysis. We can calculate the prior probability that the valve covers are in tact using formula (11.15:

$$\begin{aligned} &Pr(\mathit{valve\ covers\ ok} \mid \mathit{filter\ fire}).Pr(\mathit{filter\ fire}) = \\ &\quad Pr(\mathit{valve\ covers\ ok} \wedge \mathit{filter\ fire}) \end{aligned} \tag{11.29}$$

The following calculation introduces the conditional probabilities in Figure 11.12.

$$Pr(\mathit{valve\ covers\ ok} \wedge \mathit{filter\ fire})$$

$$= \begin{pmatrix} 1.0x0.98 & 0.98x0.02 \\ 0.0x0.98 & 0.02x0.02 \end{pmatrix} \tag{11.30}$$

$$= \begin{pmatrix} 0.98 & 0.0196 \\ 0.0 & 0.0004 \end{pmatrix} \tag{11.31}$$

In order to derive the prior probability $Pr(\mathit{valve\ covers\ ok})$ from $Pr(\mathit{valve\ covers\ ok} \wedge \mathit{filter\ fire})$ we have to use a procedure called marginalisation. This is characterised as follows:

$$Pr(A) = \sum_{B} Pr(A, B) \tag{11.32}$$

This can be applied to the matrix in (11.31) to derive $Pr(\mathit{valve\ covers\ ok}) = (0.9996, 0.0004)$. In other words the prior probability that the valve covers are in tact is just over 99%. Jensen provides more details on both the theoretical underpinning and the practical application of Bayesian networks [398]. The key point is that the underlying calculus provides investigators with a sophisticated analytical toolkit that can be used to supplement the less formal reasoning supported by the Bayesian network illustrated in Figure 11.18. The calculus can be used to derive prior and contingent probabilities depending on the nature of the information that is provided. Unfortunately, as can be seen from the preceding example, that application of these techniques can be complicated even for specialists who have considerable expertise in Bayesian analysis. For this reason, most practical applications of the approach rely upon the support of automated tools such as Hugin [399]. The previous calculations also relied upon the adaptation of models that were first developed to support medical diagnosis. This introduces the possibility that errors may have been introduced into the calculations as a result of attempting to reformulate the models to yield particular insights into the Nantcoke case study.

To summarise, the final two sections of this chapter have looked beyond the well-understood deterministic models of causation that have been embodied within incident and accident analysis techniques. The intention has been to determine whether investigators might benefit from recent developments in the theory and application of probabilistic models of causation. We have seen how this area promises many potential benefits. For example, the partition model and Baysian approaches can deal with the uncertainty that characterises the initial stages of an investigation. The importance of this should not be underestimated. *Given the increasing complexity and coupling of modern, safety-critical systems, it is inevitable that investigators will find it more and more difficult to determine a unique cause for many adverse incidents.* The Rand report into the National Transportation Safety Board (NTSB) repeatedly points to the increasing length of time that must be spent before analysts can identify the causes of many recent failures [482].

It is difficult to assess the true potential of these techniques because they have not been widely applied to support the causal analysis of adverse occurrences. In their current form there is little chance that they will be accessible to many investigators. Tool support must be provided. Methods and procedures must also be developed to help investigators learn how to apply these techniques without necessarily requiring a full understanding of the underlying theories that support the analysis. The use of Why-Because graphs as a central feature of WBA provides a useful prototype in this respect. The previous analysis has, however, identified several key issues that must be addressed before these more applied techniques will yield tangible benefits. In particular, there must be some means of assessing prior probabilities if investigators are to exploit Bayesian techniques for analysing causality through contingent probabilities. Dembski summarises this argument as follows:

> "Bayesian conceptions of probability invariably face the problem of how to assign prior probabilities. Only in special cases can prior probabilities be assigned with any degree of confidence (e.g., medical tests). so long as the priors remain suspect, so does any application of bayes' theorem. On the other hand, when the priors are well-defined, Bayes' theorem works just fine, as does the Bayesian conception of probability. To sum up then, there is no magic bullet for assigning probabilities" [200]

There may not be any general-purpose magic bullet but the previous pages have, at least, identified two potential solutions that might work as a means of assigning priors within the specialist domain of incident investigation. Firstly, we have shown how subjective probabilities can be derived using the lottery-based procedures of Von Neumann and Morgenstern [626] or of March and Simon [513]. In general these are difficult to apply because individual attitudes to risk make it difficult to interpret the expressed preferences that support inferences about subjective probabilities. We are not, however, dealing with a general population. Investigators are, typically, trained in the fundamentals of reliability and risk assessment. There is some prospect, therefore, that this method might yield better results than the more general studies of decision making under conditions of economic uncertainty.

The second, perhaps obvious, point is that we are not attempting to assign prior probabilities with complete ignorance about the nature of previous failures. In many ways, the entire purpose of an incident reporting system is to provide precise the sorts of quantitative information that is necessary in order to calculate the prior of Bayesian inference! It is paradoxical, therefore, to deny the usefulness of this data in a book that is devoted to the potential benefits of incident reporting. Unfortunately, as we have seen, we cannot trust the statistics that are extracted from national and international systems. Previous chapters have cited various estimates for the under-reporting of adverse occurrences. For instance, the Royal College of Anaesthetists estimates that only 30% of adverse medical incidents are voluntarily reported [715], Barach and Small estimate that this figure lies somewhere between 50 and 95% [66]. Chapters 5 and 15 describe techniques that can be used to assess the extent of this problem. For example, workplace monitoring can be used to identify the proportion of adverse incidents that occur within a given time period in a representative team. The results of this analysis can then be compared with incident submission rates by a similar workgroup. This is not a panacea. Even if we can assess the contribution rate within a reporting system, there is still no guarantee that we can trust the data that has been gathered about an incident. Consider the Nanticoke case study, if we wanted to gather data about the prior probability of fuel from a filter being involved in a fire, we would have to be sure that previous incidents were correctly analysed and indexed to indicate that this had indeed been a factor in previous incidents. The reliability of data about prior probabilities would be compromised if other investigators incorrectly diagnosed an incident as a filter fire when it was not. It data would also yield incorrect priors if investigators failed to diagnose this fuel source when it had contributed to an incident. Chapter 15 describes a statistical technique that can be used to identify and correct for these potential biases. For now it is sufficient to observe that this approach will only work if investigators have already performed a causal analysis of previous incidents. From this it follows that the application of Bayesian techniques may ultimately depend upon and support the use of more deterministic analysis.

11.4 Comparisons

Previous sections have reviewed a number of different techniques that can be used to support the causal analysis of safety-critical incidents. The diversity of these techniques makes it important that investigators and their managers have some means of assessing the support offered by these different approaches. Unfortunately, a range of practical, theoretical and also ethical issues complicate any attempt to perform comparative evaluations of causal analysis techniques:

- *the costs of learning new techniques.* Considerable training is required before investigators can apply some of the causal analysis techniques that we have considered. A significant level of investment would be needed to sponsor the evaluation of mathematical approaches unless investigators already had an appropriate training in the use of logic or of statistical reasoning. Similarly, it is difficult not to underestimate the problems associated with the independent application of Tier Analysis. Previous sections have emphasised the political and social pressures that affect the attribution of root causes to different levels within complex commercial organisations. Any investment in the evaluation of these techniques would carry the significant risk that they might not benefits the sponsoring organisation.

- *the costs of applying new techniques.* The investment that is required in order to train investigators to use particular analysis techniques must be supplemented to meet the costs associated with applying those techniques. This book has argued that computer-controlled automation supports increasingly complex application processes [675]. At the same time, incident investigations have been further complicated by the increasing appreciation that organisational, technical and human factors contribute to the causes of many 'failures'. These two influences have complicated the tasks associated with incident investigation. They are taking longer to complete and increasingly require the participation of multidisciplinary teams of investigators [482]. These increasing costs have not, to date, justified the allocation of resources to determine whether certain causal analysis techniques help to control the overall expenditure on incident investigations.

- *practice effects and the problems of fatigue.* Empirical test-retest procedures provide means of reducing the costs associated with the 'live' use of analysis techniques within multidisciplinary investigation teams. Investigators are presented with an example of a causal analysis technique being applied to a particular case study incident. The relative merits of that particular technique are assessed by asking investigators to answer comprehension questions, to complete attitude statements about the perceived merits of the approach and by timing investigators during these various tasks. The same procedure is then, typically, repeated for a number of further causal analysis techniques after a short break. This creates several experimental problems. For example, investigators can use the insights that were gathered from the first analysis technique to help answer questions about the second. One would, therefore, expect that the quality of the analysis might improve. On the other hand, investigators will also suffer from increasing fatigue as the evaluation proceeds. This, in turn, will impair their performance. These practice and fatigue effects can be addressed by counter-balancing. Different analysis techniques are applied in a different order by different investigators. One group might be presented with a STEP analysis and then a MORT analysis. This order would be reversed for another group. Such studies do not, however, provide any insights into the application of particular techniques over prolonger periods of time.

- *the problems of assessing learning effects.* The test-retest procedures, described in the previous paragraph, do not provide information about the long-term support that may be provided by a causal analysis technique. There studies also often yield subjective results that are strongly in favour of techniques that are similar to those which investigators are already familiar with. These potential biases create many problems. For instance, the results of a test-retest validation may simply indicate 'superficial' preferences based on a brief exposure to a relatively simple case study. These results may not be replicated if investigators actually had to apply a technique during a 'live' investigation. For example, we have described the results of an evaluation conducted using off-shore oil workers in which techniques that achieved the lowest subjective satisfaction ratings also yielded the highest comprehension and analysis scores [403]! Similarly, innovative techniques can often be undervalued if they provide significant long-term benefits that are not readily apparent during a cursory inspection.

- *the difficulty of finding 'realistic' examples.* Test-retest techniques reduce the costs associated with the validation of causal analysis techniques. The investment associated with training investigators is avoided because they, typically, are not required to apply the techniques themselves. The costs associated with applying the technique are, therefore, also avoided. Investigators are only committed to an initial assessment of existing case studies. This raises further concerns. In particular, the choice of case study may influence the investigators' responses. This is a significant issue because, as we have seen, techniques that focus on the managerial and organisational causes of failure may provide few insights into the failure of technical barriers. The test-retest procedure must, therefore, be replicated with several different case studies to provide a representative sample of the potential incidents that must be addressed. This, in turn, raises concerns that the individual preparing the case studies may also introduce potential biases that reflect their own experience in applying particular techniques. Some of these problems are addressed by accepting the costs associated with longitudinal studies of investigators applying causal analysis techniques. Given that high-consequence incidents will be rare events, even this approach provides no guarantee that investigators will meet a sufficient range of potential failures.

- *the difficulty of ensuring the participation of investigators.* Many of the previous problems relate to the difficulty of identifying an appropriate experimental procedure that can be used to support comparisons between causal analysis techniques. These issues often play a secondary role to the practical difficulties that are associated with ensuring the 'enthusiastic' participation of investigators in these studies. As we have seen, investigatory 'methodologies' are often intended to improve the *accuracy* of investigations by imposing *standard* techniques [73]. They constrain an individual's actions in response to a particular incident. It is, therefore, essential that to encourage the support and participation of investigators in the evaluation process. Any technique that under-values the existing skill and expertise of investigation teams is unlikely to be accepted. Similarly, the techniques that are being assessed must be adequately supported by necessary training material that is pitched at a level that can be understood by its potential users. Above all, the comparative evaluation of a causal analysis technique must not be misinterpreted as a comparative evaluation of incident investigators.

- *the ethical issues that stem from studying the causal analysis of incidents.* We have been involved in several studies that have performed empirical comparisons of different causal analysis techniques. These evaluations often raise a host of ethical questions for the organisations that are involved. If new techniques are introduced for a trial period then many industries require that these approaches should at least be as 'good' as existing approaches. This creates an impasse because such reassurances cannot be offered until after the evaluation has been conducted. This often forces investigators to continue to apply existing techniques at the same time as a more innovative technique is being trialed. At first sight, this replicated approach seems to offer many benefits. Investigators can compare the results that are obtained from each technique. It can, however, lead to more complex ethical issues. For instance, the application of novel causal analysis techniques can help to identify causal factors that had not previously been considered. In extreme cases, it may directly contradict the findings of the existing technique. Under such circumstances, it can be difficult to ignore the insights provided by the approach when the consequences might be to jeopardise the future safety of an application process.

The following pages build on this analysis. They provide a brief summary of several notable attempts that have been made to evaluate the utility of causal analysis techniques. As will be seen, the individuals and groups who have conducted these pioneering studies often describe them as 'first steps' or 'approximations' to more sustained validation exercises.

11.4.1 Bottom-Up Case Studies

Different causal analysis techniques offer different level of support for the analysis of different causal factors. For instance, MORT provides considerable support for the analysis of managerial and organisational failure. In contrast, early versions of this technique arguably lack necessary guidance for the technical analysis of hardware and software failures. In contrast, ECF analysis lacks any causal focus and, therefore, offers a broader scope. We have seen, however, that investigators must recruit supplementary tier analysis and non-compliance analysis to focus on particular human factors, managerial and organisational causes of an incident.

It is important to emphasise that the scope of causal analysis techniques is not static. Van Vuuren perceives a cycle in different industries [844]. A focus on individual blame and on isolated forms of equipment failure leads on to a focus on the organisational causes of incidents: This change in focus has altered the 'status quo' of safety related research and led to a number of innovative tools for causal analysis, including Tripod and PRISMA. Unfortunately, there has been a tendency for some organisations to accept that organisational failure is the end point in this process. In this Whig interpretation, accident and incident investigation has culminated in an acceptance of 'systemic' failure as the primary cause of incident investigation. Causal analysis techniques that identify the organisational precursors to systemic failures must, therefore, be chosen over those that focus more narrowly on the technical and human factors causes of incidents and accidents.

This argument raises a number of concerns. Firstly, it is unlikely that our view of incidents and accidents will remain unchanged over the next decade. The increasing development of incident reporting systems is likely to provide us with access to failure data on a scale that has never before been possible. In particular, the computer-based tools that are described in Chapter 15 already enable investigators to search through millions of reports to identify trends and causal factors that were not anticipated from the exhaustive, manual analysis of local data sources [410]. The current focus on organisational and managerial issues may, therefore, be superceded as we learn more about the causes of failure. Secondly, the focus on organisational issues is not an end in itself. We know remarkably little about the organisational and managerial causes of failure [444, 701, 839]. From this it follows that current techniques that specifically address these issues may actually fail to identify important causal factors. Indeed, many of this new generation of techniques have been attacked as premature. Researchers have pointed to particular theoretical weaknesses that are perceived to create practical problems for the investigators who must apply them:

> "The distinction between active and latent failure is the most important one in order to understand the difference in impact of different kinds of human failure. However, in his discussion Reason only focuses on the human contributions at different stages during accident causation, without providing insight into whether these human contributions result in technical, human or organisational problems. The eleven General Failure Types that are listed for Tripod are... a combination of technical, human and organisational factors, and are also a combination of causes/symptoms and root causes. For example, hardware problems are likely to be caused by incorrect design and the category organisation refers to failures that can cause problems in communication, goal setting, etc. This might be acceptable for an audit tool, however, it is not for incident analysis. Although claiming to focus on management decisions, no definition of management or organisational failure is provided. The lack of knowledge of how to model organisational failure in the area of safety related research states the importance of a bottom-up approach, using empirical incident data as a main input for new models and theories to be developed." [844]

These criticisms undervalue the pioneering role that Tripod played in re-focusing attention on the managerial and organisational factors that compromise barriers and create the context for latent failures. Van Vuuren does, however, make an important point when he urges that any evaluation of incident investigation techniques should be based on empirical data, derived from bottom-up investigations. He exploited this approach to assess the utility of the PRISMA technique. A series of case studies were conducted to demonstrate that this approach might support the causal analysis of incidents in a wide range of different domains. He also sought to validate PRISMA by applying it to different case studies within the same domain. For instance, he developed variants of the Eindhoven Classification Model to analyse incidents reported in the steel industry. He began by looking at coke production. Coke is a solid substance that remains after gases have been extracted from coal and is primarily used as fuel for blast furnaces. The company that he studied had an annual production of approximately five million tons of pig-iron. This required more than two million tons of coke from two different plants. His study focussed on one of these plants which employed 300 people in a 'traditional hierarchical organisation'. His study of fifty-two incidents revealed the distribution of causal factors illustrated in Table 11.13. The coke plant lies at the beginning of the steel making process. It provides fuel for the blast furnaces that produce pig-iron. He, therefore, conducted a second case study involving a plant that transformed pig-iron from the blast furnaces into steel. Table 11.14 presents the causal classification that was obtained for twenty-six incidents that were analysed using PRISMA in this second case study.

	Organisational	Technical	Human	Unclassifiable	Total
No. of root causes	111	67	126	13	317
Percentage	35%	21%	40%	4%	100%

Table 11.13: Distribution of root causes in Coke Production [844]

	Organisational	Technical	Human	Unclassifiable	Total
No. of root causes	73	46	57	5	181
Percentage	40%	25%	32%	3%	100%

Table 11.14: Distribution of root causes in Steel Production [844]

As mentioned, Van Vuuren was anxious to determine whether PRISMA could be successfully applied to a range of different domains. He, therefore, studied that application of the technique within both an Accident and Emergency and an Anaesthesia department. These different areas within the same healthcare organisation raised different issues in the application of a causal analysis technique. The work of the Accident and Emergency department fluctuated from hour to hour and was mainly staffed by junior doctors. In contrast, the Anaesthesia department provided well-planned and highly technical working conditions. It was mainly run by experienced anaesthetists. The insights gained from applying PRISMA within these institutions were also compared from its application in an institution for the case of the mentally ill. This institution had experienced nine incidents over a twelve month period that resulted in the death of eight of their residents and one near miss where the resident involved could barely be saved from drowning in the swimming pool at the institution. The direct causes of death varied between three cases of asphyxiation, three traffic accidents outside the main location of the institution and two drownings while taking a bath. The results of the causal analysis are summarised in Table 11.15.

	Organisational	Technical	Human	Patient related	Unclassifiable	Total
No. of root causes	29	3	24	11	4	71
Percentage	41%	4%	34%	15%	6%	100%

Table 11.15: Distribution of root causes in Mental Health Study [844]

Van Vuuren's work is important because it illustrates the use of a bottom-up approach to the validation of causal analysis techniques [844]. He provides direct, first-hand insights into the strengths and weaknesses of the PRISMA approach in a range of different application domains. This approach can be contrasted with the highly-theoretical comparisons that have been made by the proponents of other techniques. Unfortunately, the Van Vuuren's results cannot easily be applied to guide any decision between the different techniques that have been introduced in previous paragraphs. We simply lack the necessary data to make such a comparison. Techniques such as MORT have been widely applied in a range of different industries but there have been few recent attempts to systematically collate and publish the experience gained from the application of this approach. Other techniques, such as WBA and the statistical partition approaches, are relatively new and have only been validated against a small number of incidents and accidents.

Van Vuuren's approach is also limited as a basis for comparisons between causal analysis techniques. He was involved in the analysis of the case studies. It can, therefore, be difficult to assess how important his interventions were in the adoption and adaptation of the PRISMA technique. It must also be recognised that the case studies were not simply intended to provide insights into the relative utility of this approach compared to other causal analysis techniques. As can be seen, the results in Table 11.13, 11.14 and 11.15 provide no insights into how easy or difficult it was to apply PRISMA. Nor do they suggest that the findings of one investigation would be consistent with those of a previous study of the same incident. Van Vuuren was not primarily interested in the criteria that make one causal analysis technique 'better' than another. The primary motive was to learn more about the nature of organisation failure in several different industries. In contrast, Benner has applied a set of requirements to assess the utility of a wide range of investigatory methodologies.

11.4.2 Top-Down Criteria

The previous paragraphs have illustrated the diverse range of of causal analysis techniques that might be recruited to support incident investigation. This diversity is also reflected within investigatory organisations. Benner conducted a pioneering study into the practices of seventeen US Federal agencies: Consumer Product Safety Commission; Department of Agriculture; Department of the Air Force; Department of the Army; Department of Energy; Department of Labour; Mine Safety and Health Administration - Department of Labour; Occupational Safety and Health Administration (OSHA); Coast Guard; Department of Transportation; Federal Highways Administration - Department of Transportation; General Services Administration; Library of Congress; NASA; National Institute of Occupational Safety and Health; NTSB; Navy Department; Nuclear Regulatory Commission; National Materials Advisory Board - Panel on Grain Elevator Explosions [73]. He identified fourteen different accident models: the event process model, the energy flow process model, fault tree model; Haddon matrix model; all-cause model; mathematical models; abnormality models; personal models; epidemiological models; pentagon explosion model; stochastic variable model; violations model; single event and cause factors and a chain of events model. The term 'accident model' was used to refer to "the perceived nature of the accident phenomenon". Benner reports that these models were often implicit within the policies and publications of the organisations that he studied. He, therefore, had to exploit a broad range of analytical techniques to identify the investigators' and managers' views about the nature of accidents and incidents. Benner's study also revealed that these different models supported seventeen different investigation methodologies: event analysis; MORT; Fault Tree Analysis; NTSB board and inter-organisational study groups; Gannt charting; inter-organisational multidisciplinary groups; personal judgement; investigator board with intraorganisational groups; Baker police systems; epidemiological techniques; Kipling's what, when, who, where, why and how; statistical data gathering; compliance inspection; closed-end-flowcharts; find chain of events; fact-finding and legal approach; 'complete the forms'. The term 'accident methodology' refers to "the system of concepts, principles and procedures for investigating accidents" [73].

Benner's findings have a number of important consequences. He argues that the choice of accident methodology may determine an organisation's accident model. For instance, the application of the MORT technique would naturally lead to a focus on managerial issues. The use of Gannt charts would, similarly, suggest an accident model that centres on processes and events. Benner also observed the opposite effect; accident models can predispose organisations towards particular methodologies. An enthusiasm for epidemiological models leads to the development and application of an epidemiological methodology. He also identifies a third scenario in which an analysis method determines the accident model and investigation methodology but neither the model nor the investigatory methodology particularly influences each other. One interpretation of this might be situations in which organisations enthusiastically impose analytical techniques upon their investigators without considering whether those techniques are widely accepted as being consistent with the investigators' perception of an accident or incident.

A number of objections can be raised both the Benner's approach and to his analysis. For example, he used interviews to extract implicit views about models and methodologies. The findings of these meetings were supported by an analysis of documents and statutes. Previous sections in this book have enumerated the many different biases that can make it difficult to interpret these forms of evidence. Conversely, this distinction between model and methodology can become very blurred. The relatively broad definition of the term 'methodology' seems to imply that it contains elements of an accident model. The relationship between these two concepts is discussed but it is not the focus of Benner's work [73]. His investigation looks beyond the causal analysis that is the focus for this chapter, however, this work does identify a number of general problems:

> "Little guidance exists in the accident investigation field to help managers or investigators choose the best available accident models and accident investigation methodologies for their investigation... No comprehensive lists of choices, criteria for the evaluation or selection, or measures of performance (have) emerged to help accident investigators or managers choose the "best" accident model and investigative methodology." [73]

In order to address this problem, Benner proposed a series of criteria that might be used to guide a comparative evaluation of both accident models and their associated methodologies. A three point rating scheme was applied in which 0 was awarded if the model/methodology was not likely to satisfy the criterion because of some inherent shortcoming, 1 was awarded if the model/methodology could satisfy the criterion with some modification and 2 indicated that the model/methodology was likely to satisfy the criterion. Benner applied this scheme without any weightings to differentiate the relative importance of different criteria. He also acknowledges that the procedure was flawed "undoubtedly, ratings contained some author bias". The contribution of this work, arguably, rests on criteria that helped to guide his evaluation of accident models and methodologies.

The following list summarises Benner's conditions for models that reflect the perceived nature of accident phenomena. As will be seen, these criteria cannot be directly applied to assess the relative merits of causal analysis techniques. Most of the requirements support reconstructive modelling and simulation. Benner's methodological requirements have greater relevance for the content of this chapter. The model criteria are presented here, however, for the sake of completeness. This also provides an overview of Benner's more general comparison of investigatory techniques. It should be noted that we have redrafted some of the following criteria to reflect our focus on incident analysis rather than accident investigations:

1. *realistic.* This criteria focuses on the expressiveness of an incident model. Benner argues that it must capture the sequential and concurrent aspects of an adverse occurrence. It must also capture the 'risk-taking' nature of work processes.

2. *definitive.* Any model must describe the information sources that must be safe-guarded and examined in the aftermath of an incident. Ideally, the model must be composed from 'definitive descriptive building blocks' that enable investigators to set more focussed objectives during the investigatory process.

3. *satisfying.* The model must fit well with the investigatory agency's wider objectives, including any statutory obligations. It should not compromise the agencies 'credibility' or the technical quality of its work.

4. *comprehensive.* The model must capture both the initiating events and the consequences of an incident. It must capture all significant events. It must avoid ambiguity or gaps in understanding.

5. *disciplining.* The incident model must provide a rigorous framework that both directs and helpts to synchronise the activities of individual investigators. It should also provide a structure for the validation of their work.

6. *consistent.* This criterion urges that the incident model must be 'theoretically consistent' with the investigatory agencies safety program.

7. *direct.* The model must help investigators to identify corrective actions that can be applied in a prompt and effective manner. It should not be necessary to construct lengthy narrative histories before an immediate response can be coordinated.

8. *functional.* Accident models must be linked to the performance of worker tasks and to workflows. It should enable others to see how the performance of these tasks contributed to, mitigated or exacerbated the consequences of incident.

9. *noncausal.* "Models must be free of incident cause or causal factor concepts, addressing instead full descriptions of incident phenomenon, showing interactions among all parties and things rather than oversimplification; models must avoid technically unsupportable fault finding and placing of blame" [73].

10. *visible.* Models must help people to see relevant aspects of an incident. This should include interactions between individuals and systems. These representations must be accessible to investigators and to members of the general public who may themselves be 'victims' of an incident.

These criteria illustrate the way in which Benner's accident or incident models can be seen as models or templates for the incident reconstructions that have been described, for instance, in Chapters 8 and 10. The recommendation that models must capture the 'initiating events and the consequences of an incident' was a recurring theme of the earlier sections in this book. There are, however, some important differences between the approach developed in his paper and the perspective adopted in this book. For instance, Benner's criteria intend that accident models should be 'noncausal'. In contrast, we have argued that investigators cannot avoid forming initial hypotheses about the causes of an incident during the early stages of an investigation. Investigators must be encouraged to revise these working hypotheses as their work develops [850].

Benner's concept of an accident or incident model helps to determine what investigators consider to be relevant when analysing a particular 'failure'. In consequence although his model requirements focus on primary and secondary investigations, they indirectly determine the information that will be available to any causal analysis. In addition to the model criteria, list above, Benner proposes the following list of methodological requirements:

1. *encouragement.* This criteria argues that methodologies must encourage the participation of different parties affected by an investigation. Individual views must also be recognised and protected within such an approach.

2. *independence.* It is also important that methodologies should avoid 'blame'. The role of management, supervisors, employees must be recognised within the methodology.

3. *initiatives.* Personal initiatives must also be supported. It should produce evidence about previous failures that promotes intervention and shows what is needed to control future risks in the workplace.

4. *discovery.* Methodologies must support the timely discovery of information about an incident. It should also be clear when the discovery of such information may be delayed until a credible sample has been established or until "causality requirements are met" [73].

5. *competence.* This criteria argues that methodologies must leverage employees' competence. For example, it should be supported by training. This, in turn, should support the detection, diagnosis, control and mitigation of risk.

6. *standards.* Methodologies must provide credible and persuasive evidence for setting or reinforcing safety standards. It must also enable investigators to document and monitor the effectiveness of those standards over time.

7. *enforcement.* This criteria is intended to help ensure that a methodology can be used to identify potential violations. The methodology must explore deviations from expected norms. Compliance problems must be identified.

8. *regional responsibility.* In the US context, methodologies must help individual States to ensure that incident reports provide consistent and reliable accounts of accidents and incidents. More generally, methodologies must identify the role that regional organisations can play in identifying safety objectives from individual reports.

9. *accuracy.* Methodologies must validate the products of an investigatory process. It must assess the technical 'completeness, validity, logic and relevance' of these outputs.

10. *closed loop.* Methodologies must close the loop with design practices. Previous risk assessments often contain information about anticipated failure modes. These can be assessed against what actually did happen during an incident. In turn, future safety assessments can be informed by the results of previous investigations.

Benner identified the personal bias that affected his analysis. The detailed scores from his investigation are, therefore, less interesting than the criteria that drove the evaluation. In passing, it is worth noting that models which tended to represent accidents as processes were rated most highly according to the criteria listed above. These included the event process model and the energy flow process model. Elements of both of these techniques were incorporated into Benner's work on P-Theory and the STEP method mentioned previously. MORT was also ranked in the top three places according to these criteria. Similar findings were reported for the methodologies that were examined. Events analysis was rated most highly. MORT was ranked in second place assuming that it incorporated the ECF extensions described in Chapter 10 [430].

Many of the techniques that were assessed by Benner continue to have a profound impact upon existing investigatory practice. For instance, MORT is still advocated as a primary analysis technique by the US Department of Energy almost two decades after it was originally developed. Many aspects of accident and incident investigation have, however, changed since Benner first published his analysis. In particular, there has been a growing recognition of the organisational causes of adverse occurrences [701]. Software related failures also play a more significant role in many investigations [411]. The following paragraphs, therefore, use Benner's criteria to structure an evaluation of the causal analysis techniques that have been presented in previous pages.

Encouragement

This criteria was intended to ensure that methodologies encourage participation in an investigation. It is difficult, however, for many people to follow the detailed application of some causal techniques that have been presented in this chapter. This might act as a disincentive to participation during certain stages of a causal analysis. For instance, it can be hard to follow some of the statistical techniques if people are unfamiliar with their mathematical underpinnings. Similarly, the Explanatory Logic that supports WBA can be difficult to communicate to those without a training in formal logic. Fortunately, the proponents of mathematical techniques for causal analysis have recognised these objections. WBA is supported by a diagrammatic form that provides important benefits for the validation of any proof. Similarly, individuals can participate in the application of Bayesian techniques, for instance through the procedures of subjective risk assessment, without understanding all of the formulae that an investigator may employ during the more final aspects of the analysis.

These communications issues also reveal tensions within Benner's criteria. Mathematically-based techniques, typically, benefit from well-defined syntax and semantics. They provide proof rules that offer objective means of establishing the completeness and consistency of an analysis. These strengths support the accuracy criteria, assessed below, but are achieved at the expense of potentially discouraging some participation in the application of the analysis techniques. Conversely, the accessibility of Tripod, of ECF, MES and of STEP all encourage wider participation in the analysis. There is, however, a strong subjective component to the forms of analysis that are supported by these techniques. There is also a danger that participation without strong managerial control can compromise the findings of any analysis.

Independence

This criterion is intended to ensure that any methodology addresses the 'full scope' of an incident. It should consider the role of management, supervisors and employees without connotations of guilt or blame. Some techniques, including Tier analysis and MORT, provide explicit encouragement for investigators to consider these different aspects of an incident. As we have seen, however, there is a strong contrast between causal analysis techniques that offer checklist support and those that expect investigators to scope their analysis. It would be perfectly possible to apply tier analysis to an incident but for the analyst to overlook a particular sections of a management structure. In contrast, the MORT diagram provides explicit guidance on the roles that contribute to an incident or accident. The difficulty with this approach is that it can be difficult for analysts to match the details of a particular situation to the pre-defined scope in a checklist-based approach.

This criteria introduces further tensions between the Benner criteria. For example, techniques that encourage an independent assessment of the various contributions to an incident and accident can also lead to the tensions and conflict that discourage widespread participation. Many analysis techniques almost inevitably create conflict as a by-product of their application within the investigatory process. For instance, there are many reasons why non-compliance analysis creates resentment. It gains its analytical purchase from the observation that many incidents and accidents stem from individual and team-based 'failures' to follow recognised standards. Chapter 3 has also explained how observations in the health care and aviation domains have shown that operators routinely violate the myriad of minor regulations that govern their working lives. These violations do not have any apparent adverse consequences and often help individuals to optimise their behaviour to particular aspects of their working environments. In extreme examples, they may even be necessary to preserve the safety of an application process. Non-compliance analysis should reveal these violations. Investigators must, therefore, be careful to pitch their recommendations at a level that is intended to achieve the greatest safety improvements without *unnecessarily* alienating other participants in the investigatory process. Other causal analysis techniques raise similar issues. For example, Tier analysis is unlikely to promote harmony. Investigators successively associate root causes with higher and higher levels within an organisation. As mentioned, this encourages the independent analysis of the many different parties who can contribute to an adverse occurrence. However, it can lead to strong feelings of guilt, blame and anxiety as root causes are successively associated with particular levels in a management structure.

Initiatives

This criterion is intended to ensure that any accident or incident methodologies provide adequate evidence to encourage the focussed actions that address risks in a specific workplace. The previous sections in this book have not explicitly considered the means by which such recommendations for action can be derived from the products of any causal analysis. This is the central topic of Chapter 12. We have, however, considered how some analysis techniques can be used to derive particular recommendations. For instance, our analysis of PRISMA included a description of Classification/Action Matrices. These enable investigators to simply 'read-off' an associated action from a table once the causes of an incident have been determined by previous stages in the analysis. MORT offers similar support. Table 11.7 presented the 'Stage 2' analysis form proposed by the US Department of Energy. This is encourages analysis to enumerate the different ways in which an incident can be explained by the particular failures that are enumerated in the branches of a MORT diagram. The frequency with which specific items in the *why* branch are identified helps to establish priorities for action. These, in turn, help to justify the initiatives and interventions that are recommended by Benner's criteria. As we shall see in Chapter 15 similar summaries can be derived by collating the causal analysis of several incidents.

As before it is possible to highlight differences between the checklist and 'free form' approaches. Techniques, such as PRISMA, encourage a consistent approach to the recommendations and initiatives that are motivated by a causal analysis. Investigators have limited scope to alter the actions that are recommended by particular cells within a Classification/Action matrix. Conversely, less structured techniques enable investigators to tailor their response to the particular circumstances of an incident. The consequence of this is that without additional methodological support it is likely that different investigators will initiative different responses to very similar incidents.

Discovery

This criterion requires that an incident methodology should encourage a 'timely discovery process'. We have shown in previous paragraphs that there are tensions between Benner's criteria, for example encouragement can conflict with accuracy. This criteria illustrates a form of internal conflict. For example, checklist approaches are likely to provide relatively rapid insights because the items that they present can help to structure causal analysis. In contrast, techniques such as ECF analysis or Bayesian approaches to statistical causality are likely to take considerably longer given that investigators lack this guidance. In contrast, the application of 'raw' checklists is unlikely to yield entirely innovative insights. Investigators will be guided by the items that have already been identified by the author of the checklist. Free-form techniques arguably offer less constraints to the discovery process.

As mentioned, these methodology criteria were originally intended to support a comparison of accident investigation techniques. The causal analysis of safety-critical incidents creates new challenges. For instance, the statistical analysis of a body of incident reports can be used to yield new insights that might not emerge from the study of individual mishaps. The techniques that we have summarised in this chapter each pose different problems for the application of this form of discovery through data mining. For instance, the subjective nature of ECF analysis can make it very difficult to ensure that different investigatory teams will identify the same root causes for the same incident. This creates problems because any attempt to retrieve incidents within similar root causes will miss those records that have been 'miss-classified'. It will also yield reports that are potentially irrelevant from the perspective of the person issuing the query if they cannot follow the justification for a particular classification. In contrast, PRIMA's use of the Eindhoven Classification Model is intended to reduce inter-analyst variation by providing a high-level process to guide the allocation of particular categories of causal factors. Problems stem from the use of causal trees prior to the the use of the classification model. Subtle changes to the structure of the tree will have a significant impact on the number and nature of the root causes that are identified for each incident. This, in turn, can have a profound impact on the discovery of causal factors through the analysis of incident databases.

The argument in the previous paragraph assumes that causal analysis techniques have a measurable impact upon both the speed of an investigation and the likelihood that any investigation will yield new discoveries. As we shall see, some initial evaluations have shown that the investigators' background has a more profound impact upon these issues than their application of a particular technique. The discovery of particular causes can be determined by the investigator's familiarity with the nature of the corresponding more general causes. Individuals with human factors expertise are more likely to diagnose human factors causes [484].

Competence

The competence criterion requires that any methodology must help employees to increase the effectiveness of their work. This implies that they must be able to use a causal analysis technique in a cost-effective manner. Appropriate training must enable individuals to detect, diagnose, control and ameliorate potential risks. This criteria has clear implications for the more mathematical techniques that we have examined. The Explanatory Logic of WBA will only deliver the intended benefits of precision and rigour if those who apply it are correctly trained in its many different features. The partition approach to probabilistic causation provides a more pathological example of this. It is unclear precisely what aspects of the existing theories might actually be recruited to support incident investigation. The development of appropriate training courses, therefore, lies in the future. Conversely, these approaches offer means of objectively assessing the competence of individuals in the application of a causal analysis technique. Mathematical rules govern the use of statistical data and the steps of formal proofs. Tests can be devised to determine whether or not individuals understand and can apply these mathematical concepts. Such evaluations are more difficult to device for less formal approaches where the rules that govern 'correct' transformations are less clearly defined.

Techniques such as MORT and Tripod have already been widely adopted by commercial and industrial organisations [430, 701]. Training courses and commercial software can also be recruited to improve employee competence in the application of these techniques. PRISMA arguably rests halfway between these more commercial techniques and the more novel mathematical approaches to causal analysis. As we have seen, there is limited evidence that this approach can be used in a range of different contexts within several industries. These is, as yet, relatively little guidance on how investigators might be trained to exploit this approach. It is important to emphasise that the lack of training materials does not represent a fundamental objection to any of the techniques that have been considered in this book. Our experience in training investigators to conduct various forms of causal analysis has shown that most organisations tend to develop their own training materials. It is important that any causal analysis technique supports both their organisational priorities and also the reporting obligations that are imposed on them by other statutory and regulatory bodies.

Standards

The standards criterion requires that incident methodologies must enable investigators to identify potential flaws in their work. They must also provide a comprehensive, credible, persuasive basis for the advocacy of appropriate interventions. This criterion served as a prerequisite for the inclusion of causal analysis techniques in this book. It is possible, however, to identify a number of distinct approaches that are intended to improve the standard of incident investigation within the different approaches that we have analysed.

Arguably the weakest claim that is made for causal analysis techniques is that they provide intermediate representations, typically figures and graphs, that can be exposed to peer review during the investigatory process. ECF charting, non-compliance tables, MES flowchart all help to explicitly represent stages of analysis. This can help to identify potential contradictions or inconsistencies that might otherwise have been masked by the implicit assumptions that are often made by different members of an investigatory team.

Many of the techniques that we have studied also provide particular heuristics or rules of thumb that provide a basis for more complex forms of analysis. MES, STEP, ECF, MORT, WBA all exploit variants of counterfactual reasoning. This approach offers considerable benefits not simply because it encourages a consistent approach to the identification of causal factors. Counterfactual reasoning also provides investigators with a common means of identifying potential counter-measures. Recall that the counterfactual question can be expressed as 'A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either'. We might, therefore, consider preventing an incident, B, by devising means of avoiding B. As we have seen, however, causal asymmetry implies that investigators must be circumspect in exploiting techniques which advocate this style of analysis. Further questions arise both from the cognitive problems of reliably applying counterfactual reasoning [124] and from the practical problems of validating hypothetical reasoning about the failure of barriers, described in Chapter 10.

This criteria also urges that causal analysis techniques should be assessed to determine whether they provide a comprehensive, credible, persuasive basis for the advocacy of appropriate interventions. The interpretation of a 'comprehensive' approach depends upon the scope of the techniques. This was addressed in the previous section. It s more difficult to assess the credibility and persuasiveness of an approach. We are unaware of any studies that have directly addressed this issue as part of an evaluation of causal analysis techniques. Similar studies in other domains have, however, indicated that the application of a particular method may be less important than the identity of the individual or group *who* apply the technique [278].

Enforcement

Benner's criteria require that an incident methodology should reveal expectations about the norms of behaviour. This, in turn, helps investigators to identify non-compliance. It should also help to identify instances in which the enforcement of standards has been insufficient to prevent violations. As with the other criteria, each of the techniques that we have assessed can be argued to offer different levels of support for these aspects of 'enforcement'. For example, non-compliance analysis was integrated into our use of ECF in Chapter 10. This technique is deliberately intended to identify violations and to expose deviations from expected norms. The Explanatory Logic that supports the formal components of WBA also includes deontic operators that explicitly capture notions of obligation and permission. These have been used to represent and reason about the particular consequences of non-compliance with standards and procedures [118, 469].

A number of caveats can be made about this criterion and its application to causal analysis techniques. Firstly, it can be difficult to distinguish a willful violation from a slip or a lapse. The distinction often depends upon the analyst's ability to identify the intention of an individual or group. None of the causal analysis techniques that we have investigated support the inference of intention from observations of behaviour. The PARDIA components of WBA provide a possible exception to this criticism. Unfortunately, there are relatively few examples of this technique being used to analyse complex, operator behaviours. It remains to be seen whether this approach might be used to enable analysts to reason about the detailed cognitive causes of violation and non-compliance.

A second caveat to Benner's enforcement criteria is that numerous practical difficulties frustrate attempts to chart the differences that exist between everyday working practices and the recommendations of standards and regulations. Chapter 10 showed that it was extremely difficult for executives, managers and supervisors to keep track of the dozens of procedures and guidelines that had been drafted to guide the development of the Mars Polar Lander and Climate Orbiter projects. Previous paragraphs have also noted the high-frequency of apparently minor violations that have been noted as characteristic of expert performance within particular domains, especially aviation [672]. The 'enforcement' criterion, therefore, represents a class of requirements that are currently not adequately met by any causal analysis techniques. These criteria can be contrasted with other requirements, such as the need to provide 'standards' for causal analysis, which are arguably satisfied by all of the techniques that we have examined.

Regional responsibility

This criterion was initially drafted to ensure that individual States are encouraged to use a methodology and to take responsibility for its application within the context of U.S. Health and Safety "mandates" [73]. In contrast, we argue that causal analysis techniques must consider the different roles and objectives that distinguish regional organisations from the local groups that, typically, implement reporting systems. Some causal analysis techniques seem to be better suited to regional organisations. For instance, Chapter 10 shows how Tier Analysis associates root causes with higher levels in an organisational hierarchy. This process is likely to create conflicts between local investigators and senior members of a management organisation. Regional investigators are more likely to possess the competence and independence that are necessary to resist the pressures created by such conflicts. Other techniques, such as WBA, can be so costly in terms of the time and skills that are required to exploit them that regional and national groups must be involved in their application. In contrast, the methods that might be derived from probabilistic models of causality are likely to benefit from the information contained in large-scale datasets. Regional organisations may, therefore, be best placed to offer advice and support in the application of these techniques.

The aims and objectives of national and regional organisations are likely to be quite different from those of the local teams who are responsible for conducting an incident investigation. Regional organisations are more concerned to derive a coherent overview from a collection of incident reports than they are to understand the detailed causal factors that led to one out of several thousand or hundreds of thousands of incidents [444]. An immediate concern to mitigate the local effects of an incident are part of a wider concern to ensure that lessons are propagated throughout an industry. It is, therefore, important that regional organisations should understand and approve of the causal analysis techniques that are used to provide data for these aggregate data sources. They may also impose training requirements to ensure competence in the application of those techniques [423].

A number of further complications can, however, frustrate regional and national initiatives to exploit the products of causal analysis techniques. For instance, regional bodies are often anxious to ensure that investigators exploit similar techniques so that accurate comparisons can be made between individual reports from different areas of their jurisdiction. It is likely, however, that there will be pronounced local distortions even if different geographical regions all share the same causal analysis techniques. A succession of similar incidents or an accident with notably severe consequences can sensitise investigators to certain causes. This effect may be more pronounced for those who are most closely associated with previous incidents [410]. In consequence, very different results can be obtained when the same incidents are reclassified by investigators from different regions and even from different teams. These issues complicate attempts to share data across European boundaries in the aviation domain [423] and across State boundaries within US healthcare [453]. They are also largely orthogonal to the problems of identifying appropriate causal analysis techniques.

Accuracy

This criteria is similar to aspects of the 'standards' requirement that was introduced in the previous paragraphs. Accuracy is intended to ensure that incident methodologies can be tested for completeness, consistency and relevance. All three of these concepts are relevant to the causal analysis of safety-critical incidents. The first two directly motivated the application of formal proof techniques to support semi-formal argumentation in WBA. As mentioned, however, it can be more difficult to validate techniques that exploit precise mathematical concepts of consistency and completeness. A further caveat is that the use of formal techniques does not guarantee an error-free analysis [21]. It does, however, provide objective rules for identifying and rectifying these problems.

The other techniques that we have presented, such as STEP, MES and MORT, provide fewer rules that might be applied to assess the accuracy of a causal analysis. Instead, as mentioned, they rely upon diagrams and tables that are open for peer review. The less formal processes involved in achieving group consensus are intended to provide greater confidence than the formal manipulations of abstractions whose precise interpretation can defy even skilled mathematicians. A similar debate between informal, semi-formal and formal methods has dominated areas of computing science research for several decades [32]. The detailed comparison of the strengths and weaknesses of these different approaches lies outside the scope of this book. In particular, such comparisons have little value unless they can be substantiated by detailed empirical evidence. Later sections will briefly summarise the preliminary results that have been obtained in this area [529, 552]. Unfortunately, these findings provide limited insights into the application of particular approaches. They do not support more general conclusions about comparative benefits and weaknesses. In consequence, investigators must make their own subjective assessments of the claims that proponents make for the 'accuracy' of their causal analysis techniques.

Closed Loop

This final criterion requires that incident methodologies should be tightly integrated into other aspects of systems design and implementation. The data from incident reporting systems should inform future risk assessments. Information from past risk assessments, or more precisely the arguments that support those assessments, can also help to guide a causal analysis providing that it does not prejudice investigators' hypotheses. We have not suggested how any of the causal analysis techniques that we have examined might support satisfy such requirements. There are, however, many similarities between non-deterministic causal models and the techniques that support quantitative approaches to reliability and risk assessment. The prior probabilities of Baysian analysis can be derived from the estimates that are embodied in risk assessments, especially when reliable data cannot be derived directly from an incident reporting system.

Previous paragraphs have, however, provided an example of a risk assessment technique being used to guide the causal analysis of an adverse incident. Chapter 10 described how NASA's mishap investigation board identified a problem in the software that was designed to automatically re-establish communications links if the up-link was lost during the Polar Lander's Entry, Descent and Landing phase. This bug was not detected before launch or during the cruise phase of the flight. A Fault Tree analysis did, however, identify this possible failure mode after the Polar Lander had been lost.

Chapter 12 will return to this issue in greater detail. The relationship between risk assessment, design and incident reporting is often only considered as an after-thought by many organisations. In consequence, subjective assessments of component and system reliability are often exploited by one group within a company while others in the same organisation continue to collate data about the actual performance of those systems [414]. More generally, this same situation can arise when design groups are unwilling to approach other organisations within the same industry who have previous experience in the incidents that can arise from the operation of particular application processes. In consequence, development resources are often allocated to perceived hazards that cannot be justified by data about previous failures.

The previous paragraphs have shown how Benner's methodological criteria can be used to structure a comparison of causal analysis techniques. Some of the criteria, such as 'accuracy' and 'standards', are relevant objectives for all of the approaches that we have examined. Other requirements, such as 'encouragement', are less important for particular techniques. WBA and techniques derived from partition models of probabilistic causality focus more on 'accuracy' and 'competence'. The main weakness with this approach is that Benner fails to provide any objective procedures that might be used to determine whether or not a particular methodology satisfies a particular criterion. It is, therefore, possible for others to argue against our analysis. For example, it might be suggested that partition methods can encourage greater participation. Certainly, the diagrammatic forms of WBA do help to increase access to some aspects of this technique. There have, however, been no empirical studies to investigate the communications issues that might complicate the use of these formal and semi-formal techniques within the same approach. The following sections, therefore, briefly describe the limited number of studies that have been conducted in this area. These studies do not provide a firm basis for the comparative evaluation of causal analysis techniques. They focus on a limited subset of the available approaches. They also concentrate on incidents within particular industries. These studies do, however, illustrate the manner in which empirical evidence might be recruited to support assertions about the relative merits of these techniques.

11.4.3 Experiments into Domain Experts' Subjective Responses

Both Van Vuuren's bottom-up analysis of the PRISMA approach and Benner's application of model and methodology criteria were driven by the direct involvement of the individuals who were responsible for conducting the tests. Van Vuuren participated in the analysis that is summarised in Tables 11.13, 11.14 and 11.15. Benner performed the ratings that were derived from the lists of criteria presented in the previous section. This level of personal involvement in the validation of causal analysis techniques should not be surprising. Previous sections have summarised the practical, theoretical and ethical issues that complicate the evaluation of different causal analysis techniques. Many researchers, therefore, side-step the problems of investigator training and recruitment by conducting subjective studies based on their own application of alternative techniques. In contrast, Munson builds on the work on Benner [73] and Van Vuuren [844] by recruiting a panel of experts to validate his application of causal analysis techniques [552]. Munson began by applying a number of analysis techniques to examine a canyon fire that had previously been investigated by the U.S. Fire Service. In particular, he applied Fault Tree analysis, STEP and Barrier analysis to assess the causal factors that contributed to this incident. He then recruited his panel by choosing wildland firefighting experts rather than 'professional' investigators; this "most accurately emulates real world situations where investigators may have some investigative experience but their primary occupation and training is not in these techniques" [552]. None of the evaluators had any prior experience with accident analysis techniques. This helped to avoid any preference for, or experience of, existing approaches. Each member of the panel had a minimum of fifteen years experience in wildland fire suppression and were qualified to 'Strike Team Leader' level. Individuals were selected on a 'first come' basis.

Munson acknowledges that this may have introduced certain biases, however, he endeavoured to ensure that none of the panel consulted each other about their ratings. He was also aware that the panel members may have held preconceived ideas about the causes of the case study; "since they were evaluating the investigation methods and not the reinvestigation of the fire, bias should have been reduced" [552].

The members of the panel were asked to compare Munson's Fault Tree analysis, STEP analysis and Barrier analysis of the canyon fire by rating each technique against a number of criteria. These requirements were based on a subset of the Benner criteria [73]. As can be seen, some of these requirements apply more to reconstruction and modelling than they do to causal analysis. This can be justified by the mutual dependencies that we have stressed in previous chapters. Munson's criteria can be summarised as follows:

1. *Realistic.* Any analysis must capture the sequential, concurrent, and interactive nature of the flow of events over time.

2. *Comprehensive.* Any analysis must identify the beginning and the end of an accident sequence and there must not be any gaps in the investigator's understanding of an incident.

3. *Systematic.* Any analysis must be supported by a logical and disciplined method that allows for peer review and mutual support by all of the members of an investigation team.

4. *Consistent.* The method must be consistent and it should be possible for investigators to verify that any conclusions are correct from the information that is available.

5. *Visible.* Any analysis must discover and present the events and interactions throughout an accident sequence so that colleagues can understand the manner in which they contribute to an incident.

6. *Easy to learn.* It should be possible for investigators to learn how to apply a technique by attending a one week course. This criterion reflects Munson's focus on the fire fighting community and he acknowledges that it should not be considered when attempting to assess the 'best' analysis technique.

The experts were asked to use a ranking system that was similar to that described in the previous section; "The rating scale follows Benner's [73] approach in that until a more comprehensive scale is developed to better differentiate levels of compliance to the criterion, a more simple direct measurement scale is appropriate" [552]. For each model, they were asked to rate each criterion. A score of zero was used to denote that they did not believe that the approach met this criterion. A score of one was used to denote indicate that they believed that the approach addressed the criteria but not completely and improvement would be required. A score of two was to be awarded if the analysis technique satisfied the criterion. No weighings were applied to the results of this process because no criterion was perceived to have more significance than any other. The results from summing the individual scores showed that STEP received the highest rating; 52 from a possible 60 (87%). Fault Tree Analysis received 51 out of 60 (85%). Barrier Analysis obtained 42 out of 60 (70%). STEP was rated as the most 'comprehensive' (100%) and most 'consistent' (100%). Both Fault Tree Analysis and STEP were rated as the 'easiest to use' (90%). Barrier analysis was rated the most 'realistic' technique (90%). Fault Tree Analysis was rated as the most 'systematic' method (100It was also the most visible (90%). Two evaluators rated it as the best overall approach. Two rated STEP the highest. One assigned equal value to both STEP and Fault Tree Analysis. Barrier Analysis was not assigned the highest rating by any of the evaluators.

Munson also analysed his data to assess the level of agreement between his panel of assessors. Multivariate techniques were not used; "the number of evaluators and criteria were considered too small and would not constitute any meaningful insight" [552]. Instead, Perreault and Leigh's [674] index was used to assess inter-rater reliability. Indexes above 0.85 are considered to indicate a high degree of consensus. Levels below 0.80 require further analysis. Munson provides the following equation for the reliability index. F is the frequency of agreements between the evaluators, N is the total number of judgements and k is the number of categories:

$$I_r = [(F/N) - 1/k)][k/k - 1)]^{0.5} \tag{11.33}$$

The panel's evaluation of the six criteria for the STEP method yielded an index of 0.84 [674]. Fault Tree Analysis received 0.86 over all of the criteria. Barrier Analysis achieved a reliability index of 0.79. The inter-rater reliability for all methods was 0.84. As can be seen, only the Fault Tree assessment indicated a high degree of consensus but all other measures fell into the acceptable region identified by Perreault and Leigh [674]. If we look at levels of agreement about individual criteria it is possible to see some interesting patterns. For example, there was little agreement about whether or not Fault Tree analysis was a 'realistic' technique (0.63). STEP received the highest rating for 'comprehensiveness' and achieved an index of 1.0 for inter-rater reliability. Fault Tree analysis and Barrier analysis achieved a similar level of consensus but at a lower over rating about the 'comprehensiveness' of the techniques. Munson provides a more sustained analysis of these results [552].

The experts were each asked to provide additional comments about the applicability of each method. Munson cites a number of the responses that were provided. Ironically some of these comments reveal the experts' lack of understanding about the technical underpinning of the methods that they were asked to evaluate. The attitudes to Fault Tree analysis are particularly revealing in this context, given the key role that they play within many areas of reliability and risk assessment:

> "One evaluator liked the way Fault Tree Analysis visually presented complex events and the way it showed accidents as a chain-of-events as opposed to a single random occurrence... They thought that this method might be better at uncovering managerial/administrative latent factors contributing to the incident than the other two methods. In contrast, one evaluator responded that the STEP method appeared more stringent in revealing underlying human causal factors. They commented that STEP (and Control/Barriers Analysis) provided an approach that was more likely to distinguish more abstract human factors from hard factual data considerations and therefore be better at raising questions into human error causes... All evaluators expressed concern that Control/Barriers Analysis was inadequate in determining causal factors when applied to wildland firefighting. It had strengths in identifying needed and/or compromised barriers at an administrative level but the dynamic and highly variable aspect of the firefighting environment made its application to investigations inadequate" [552].

Munson concludes that STEP is the most 'desirable' method for the investigation of wildland firefighter entrapments. The small differences between the scores for this technique and for Fault Tree analysis suggest, however, that there are unlikely to be strong differences between these two techniques. Both were rated more highly than Barrier Analysis.

A number of questions can be raised both about the methods that Munson used and about the results that he obtained from them. Firstly, Munson was not qualified in accident or incident investigation when he undertook this study. The manner in which he applied the three techniques need not, therefore, have reflected the manner in which they might have supported an active investigation by trained personnel. Secondly, a number of caveats and criticisms have been made about his application of particular techniques. For example, Fault Tree analysis of the canyon fire breaks some of the syntactic conventioned that are normally associated with this approach, see Chapter 10. Paradoxically, these differences aid the presentation of Munson's analysis. They also make it difficult to be sure that the results from this study could be extended to the more general class of Fault Trees that obey these syntactic conventions. Thirdly, this study focuses on experts who only represent a very small cross-section of the community who are involved in accident and incident investigations. This is a universal weakness shared by all previous validation studies that we have encountered. Chapter 4 has shown that incident and accident reporting systems involve individual workers, supervisors, investigators, safety managers, regulators and so on. Benner's original 'encouragement' criteria captures some aspects of this diversity. However, experimental validations focus on the utility of causal analysis techniques for investigators or, as in this case, domain experts. Regulators might take a very different view. Fourthly, a number of minor caveats can be raised about the choice of statistical techniques that were used to analyse the data from this study. Multivariate analysis might have been applied more systematically. This could have yielded results that are easier to interpret than the piecemeal application of Perreault and Leigh's index. Finally, Munson's study specifically addresses the fire fighting domain. Several of the criteria were specifically tailored to reflect the working and training practices of this application area. Further studies are required to replicate this work in other domains.

It is important to balance these criticisms and caveats against the major contribution that has been made by Munson's work. The opening paragraphs of this section reviewed the many pragmatic, theoretical and ethical barriers that complicate research in this area. His approach successfully addresses many of these potential problems. Muson shows that it is possible to provide further evidence to support Benner's subjective analysis.

11.4.4 Experimental Applications of Causal Analysis Techniques

Previous sections have described a number of different approaches to the validation and comparative evaluation of causal analysis techniques. Van Vuuren adopted a bottom-up approach by applying PRISMA to support a number of incident reporting systems within particular industries. Benner adopted a much more top-down approach when he developed and applied a set of criteria in a subjective evaluation of accident models and methodologies. Munson used this approach as the foundation for an expert panel's evaluation of causal analysis techniques for fire fighting incidents. A limitation of Benner's approach is that it was based upon the subjective analysis of the researcher. Munson avoided this by recruiting a panel of experts. They did not, however, apply any of the methods and only provided subjective ratings based on a case study that was developed by Munson himself. Van Vuuren's study involved investigators in the application of the PRISMA technique. He, however, played a prominent role in coaching the use of this approach; "guidance was necessary to pinpoint these mistakes or lapses and by doing this to improve the quality of the causal trees and stimulate the learning process regarding how to build causal trees" [844]. This intervention was entirely necessary given the ethical issues that surround the validation of incident investigation techniques using 'live' data. The closing sections of this chapter describe an alternative approach. McElroy attempted to integrate elements of Munson's more controlled experimental technique and Van Vuuren's concern to involve potential end-users in the application of particular approaches [529].

McElroy's evaluation began with a sustained application of the PRISMA technique. He used a variant of the Eindhoven Classification Model to identify the root causes of more than one hundred aviation incidents from the ASRS dataset. This yielded approximately 320 root causes; the majority of which related to human factors issues. In order to validate his results, he recruited a number of experts to repeat his analysis of selected incidents from the study. The intention was then to compare the causal trees that they produced and the resulting root cause classification with McElroy's findings from the initial analysis. He rejected Munson's approach of recruiting domain experts, such as pilots or air traffic controllers. This was partly motivate by pragmatic reasons, such as the difficulty of securing access to participants for the study. It was also motivated by the difficulties that Munson and Benner had foreseen in training domain experts to apply novel analysis techniques, rather than simply requiring them to comment on the use of the approach be someone else. In contrast, McElroy recruited participants who had specific expertise or training in incident and accident analysis. This approach also raised problems; he found it difficult to secure the involvement of participants with similar expertise and training. Both of these factors are significant given Lekberg's results, which show that the investigator's training will influence their causal analysis of safety-critical incidents [484]. In the end he was only able to assess the application of the technique by two participants. In consequence, his findings cannot be used to support general conclusions about the PRISMA technique. They do, however, provide a glimpse of some of the individual differences that might affect the application of causal analysis techniques by incident investigators.

As mentioned, McElroy provided his participants with short synopses of incidents that had previously been submitted to the ASRS. The following paragraph provides a brief extract from the summary that McElroy used in his study:

> ACCESSION NUMBER : 412640
> DATE OF OCCURRENCE : 9808
> NARRATIVE : DEPARTING NEWPORT ARPT, AT THE TIME OF DEP, THE W HALF OF THE ARPT WAS STARTING TO FOG IN. I HOVER-TAXIED TO THE FAR E END OF THE ARPT AND WAS ABLE TO TAKE OFF IN BLUE SKIES AND UNLIMITED VISIBILITY. THIS ARPT IS SET UP FOR A CTL ZONE WHEN THE VISIBILITY IS LESS THAN 3 MI AND A 1000 FT CEILING. THERE WAS ANOTHER HELI IN THE PATTERN WHOM I WAS IN RADIO CONTACT WITH. HE GAVE ME PERMISSION TO TAKE OFF FIRST AND THEN HE WENT IN AND LANDED. ALL OF THIS WAS DONE VFR ON THE E END OF THE FIELD WHILE THE W END WAS FOGGED IN. THE STANDARD FOR THE OTHER ARPTS WITH CTL TWRS HAS BEEN IF I WAS INSIDE OF THEIR CTL ZONE AND IT WAS IN EFFECT, THEY HAVE ALLOWED ME TO WORK INSIDE THE CTL ZONE WITHOUT A SPECIAL VFR IF I WAS IN THE STANDARD VFR CONDITIONS. ALL I NEEDED TO DO WAS MAKE RPTS OF MY LOCATIONS WHILE WORKING IN THEIR AIRSPACE. AS LONG AS I WAS VFR, I DID NOT NEED A SPECIAL VFR TO BE INSIDE THE AIRSPACE. MY POINT TO ALL OF THIS IS THAT IT IS NOT TAUGHT TO NEW STUDENTS THIS WAY SO IT BECOMES MORE LIKE JUST A STORY WHEN AN OLDER PLT DOES SOMETHING LIKE THIS. IT IS LEGAL TO DO BUT NOT GOOD FOR STUDENTS TO SEE. NOT SURE OF HOW OR WHERE TO MAKE A POINT OF THIS, OR IF MAYBE IT IS NOT A RELATIVE POINT TO MAKE AT ALL. HOPE THIS IS NOT TOO CONFUSING, AND THANK YOU FOR YOUR TIME.

The first participant produced the Causal Tree shown in Figure 11.19. McElroy's study focussed more on the application of this diagram to support PRISMA's root cause analysis. A number of insights can, however, be derived from this initial stage of his evaluation. The tree took several hours to construct but, as can be seen, it is essentially a sketch of the incident. It includes inferences and judgements that cannot directly be supported from the synopsis. For instance, one note is annotated to denote that the helicopter pilot took off illegally, happy he was on a visual flight rule. Nowhere does the report state that the pilot was 'happy' with the state of affairs. Similarly, the causal tree refers to the maneuver as 'illegal' although the pilot believes that there actions were 'legal' within the control zone of the airport tower. This ambiguity reflects a lack of contextual information and the participants' limited domain knowledge. It was not, however, addressed in McElroy's analysis. A key point here is that although this evaluation ran for several hours, the participants never had the opportunity to move beyond this high-level sketching to the more meticulous analysis that would be need to demonstrate the sufficient and correctness of a causal 'explanation'. One might, therefore, infer that such an experimental procedure would have to be significantly revised if it were to be used to assess the utility of one of the mathematical techniques that we have described.

As mentioned, McElroy's aim was to determine whether participants who were training in incident analysis would confirm his own application of PRISMA. The first participant was, therefore, asked to use their diagram in Figure 11.19 to drive the categorisations of root causes using a variant of the Eindhoven Classification Model. They identified the following list of potential causes:

- Operating procedures. This is represented by the node labelled OP in the Eindhoven Classification Model of Figure 11.9. The participant identified that the incident was the result of inadequate procedures.

- Management priorities. This is represented by MP in the Eindhoven Classification Model. The participant identified that top or middle management placed pressure on the operator to deviate from recommended or accepted practice.

- Permit. This is represented by HR2 in the Eindhoven Classification Model. The participant identified that the operator failed to obtain a permit or licence for activities where extra risk was involved.

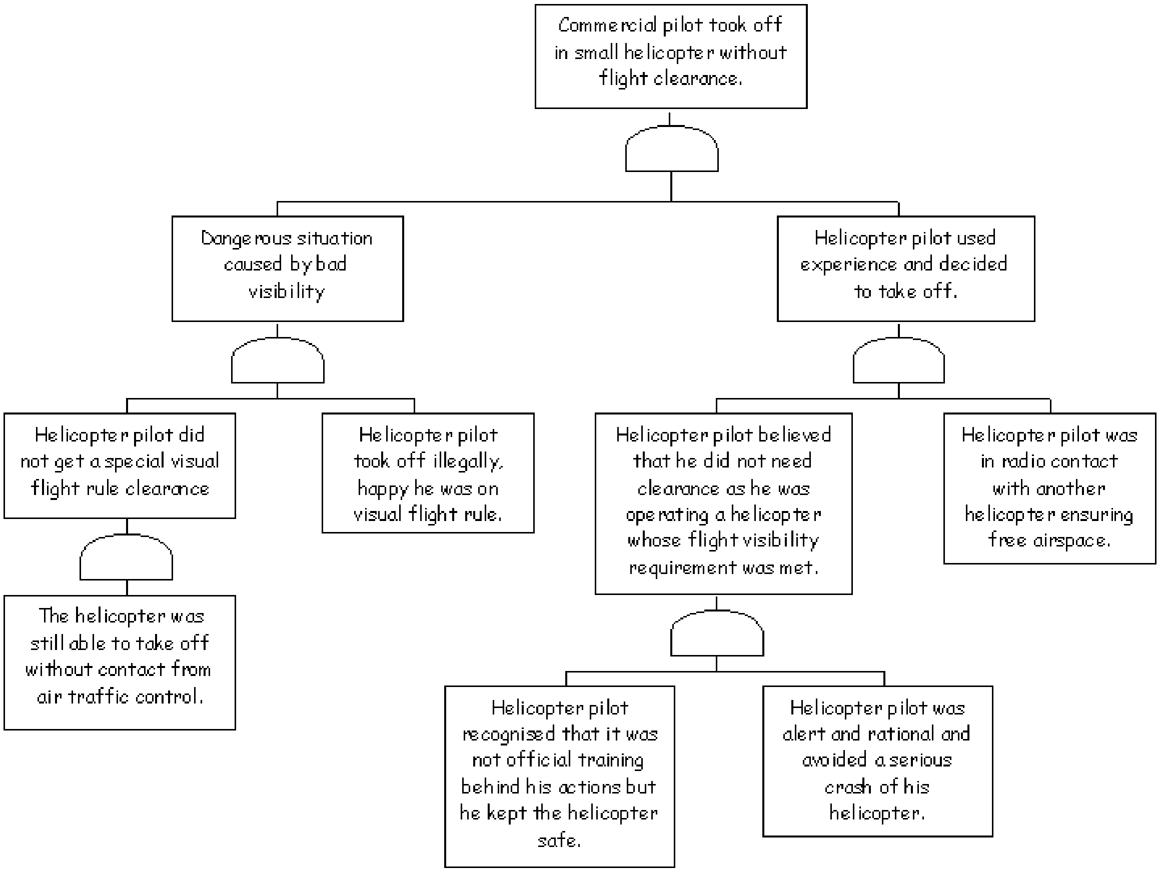
Commercial pilot took off in small helicopter without flight clearance.
Dangerous situation caused by bad visibility
Helicopter pilot used experience and decided to take off.
Helicopter pilot did not get a special visual flight rule clearance
Helicopter pilot took off illegally, happy he was on visual flight rule.
Helicopter pilot believed that he did not need clearance as he was operating a helicopter whose flight visibility requirement was met.
Helicopter pilot was in radio contact with another helicopter ensuring free airspace.
The helicopter was still able to take off without contact from air traffic control.
Helicopter pilot recognised that it was not official training behind his actions but he kept the helicopter safe.
Helicopter pilot was alert and rational and avoided a serious crash of his helicopter.

Figure 11.19: Causal Tree from McElroy's Evaluation of PRIMA (1)

- Planning. This is shown as HR5 in the Eindhoven Classification Model. The participant identified that the activity was not fully planned. Appropriate methods were not identified nor were they carried out in a well-defined manner.

- Unclassifiable behaviour. This is shown as X in the Eindhoven Classification Model. The participant also identified that some of the causal factors denoted in Figure 11.19 could not be classified using the model.

In contrast, McElroy's analysis only identified management priorities and planning as causal factors in this incident. The other three causes identified by the participant were not identified in the initial analysis. In addition, McElroy's analysis identified Goal? (HK2) as a potential cause that was not recognised by the first participant. This root cause categorisation denotes that the operator failed to identify appropriate goals or priorities for their actions. This comparison raises several issues. Firstly, the study tells us what categories the participant felt were important to the causes of the case study. It does not tell us *why* they believed them to be important. This is important because both McElroy and the first participant identified planning as a causal factor, it is entirely possible however that they had entirely different reasons for making this categorisation. Conversely, we do not know the reasons why they differed over specific elements in their causal analysis. Secondly, it is difficult to determine the justification for some of the reported conclusions made by both McElroy and the participant. Although the previous quotation is an abbreviated from of the incident report that was supplied during the study, there is no explicit indication that management priorities had caused the pilot to behave in the manner that they reported. This illustrates the more general point that we have made repeatedly in this book; it is not sufficient simply to present a causal analysis without providing a detailed justification of the reasons supporting that analysis.

As mentioned, the second evaluation focussed on an incident from the ASRS' air traffic dataset:

> ACCESSION NUMBER : 425120
> DATE OF OCCURRENCE : 9901
> NARRATIVE : WX WAS SUNNY BUT COLD, A DAY OR 2 AFTER A SNOW/ICE STORM. SABRELINER WAS TAXIING OUT FOR IFR DEP. ATC OBSERVED THE FUSELAGE WAS COVERED WITH SNOW AND ICE. ATC ADVISED THE PLT 'IT APPEARS THERE'S A LARGE AMOUNT OF SNOW AND ICE ON THE TOP OF YOUR ACFT.' THE PLT STATED 'IT'S NOT A LOT, IT'S A LITTLE, AND IT WILL BLOW OFF WHEN WE DEPART.' ON TKOF ROLL, ICE WAS OBSERVED PEELING OFF THE FUSELAGE. THIS CONTINUED AS THE ACFT CLBED OUT. ICE WAS OBSERVED FALLING ON OR NEAR A HWY JUST OFF THE DEP END OF THE RWY. THE ACFT WAS SWITCHED TO DEP, BUT A FEW MINS LATER RETURNED FOR LNDG. AS THE ACFT TAXIED IN, SIGNIFICANT ICE FORMATION WAS OBSERVED ON THE ELEVATORS. THE ACFT TAXIED TO AN FBO AND WAS DEICED BEFORE TAXIING BACK OUT FOR DEP. I SPOKE WITH THE FBO LATER. THEY SAID THEY HAD SEEN THE PLT CLRING SNOW AND ICE OFF THE ACFT BEFORE HE FIRST DEPARTED. HOWEVER, THE UPPER SURFACE OF THE ELEVATORS WAS TOO HIGH FOR THE PLT TO SEE FROM THE GND.

The second participant produced the Causal Tree shown in Figure 11.20 for this incident report. McElroy's again analysis focussed on the causal factors that were identified using a variant of PRISMA's Eindhoven Classification Model. As before, however, this diagram yields several insights into the assessment of causal analysis techniques. There is a far greater level of detail in this tree than in Figure 11.19. There is insufficient evidence to determine whether this is an artifact of individual differences between the participants or whether it stems from differences in the two incidents that they studies. As with many aspects of McElroy's work, it provides tantalising hints of deeper issues. He did not counter-balance the study so that each participant was asked to analyse each incident. This had been an intention behind the study but he ran out of time. Rather than rush the participants to perform a partial study of two incidents, he chose to allow them more time with a single incident.

Both Figures 11.19 and 11.20 are sketches. They record the participants' initial thoughts about the incidents. They follow the high-level structure proposed by the causal tree approach; left branches represent the 'failure' side while the right branch denotes 'recovery' events. There are also examples in both trees where that participants either deliberately neglect the syntax of there trees or else they did not follow the syntactic rules that were presented. In Figure 11.19, there is a minor violation with an AND gate that includes a single event. It can be argued that this represents a stylistic issue rather than a violation f any explicit syntactic rule. In this case, however, it is uncertain how to interpret the relationship between Helicopter pilot did not get a special visual flight rule clearance and The helicopter was still able to take off without contact from air traffic control. Figure 11.20 raises more questions. No checklist or protocol is linked to two events without any intervening gate. The event labelled Pilot dismiss ATC concerns is provided as an input to two different AND gates. Such techniques break the independence assumptions that are important for the analysis of more 'conventional' fault trees. These rules were, almost certainly, not presented to the participants in McElroy's study. Such annotations are, therefore, of considerable interest because they illustrate ways in which users are shaping the notation to represent the course of an incident. In future studies, it would be important to know what was intended by the event labelled No checklist or protocol. This would enable us to determine whether the notation fails to support a necessary feature or whether the training failed to convey significant syntactic constructs to the participants. Given that participants were unlikely to derive a reliability assessment from Figure 11.20 it can be argued that the independence assumption has not value for the practical application of causal trees?

As with the first participant for the helicopter case study, the second participant was also asked to use their causal tree to drive the categorisation process that is supported by the Eindhoven Classification Model. The following list summarises the categories of root causes that were identified

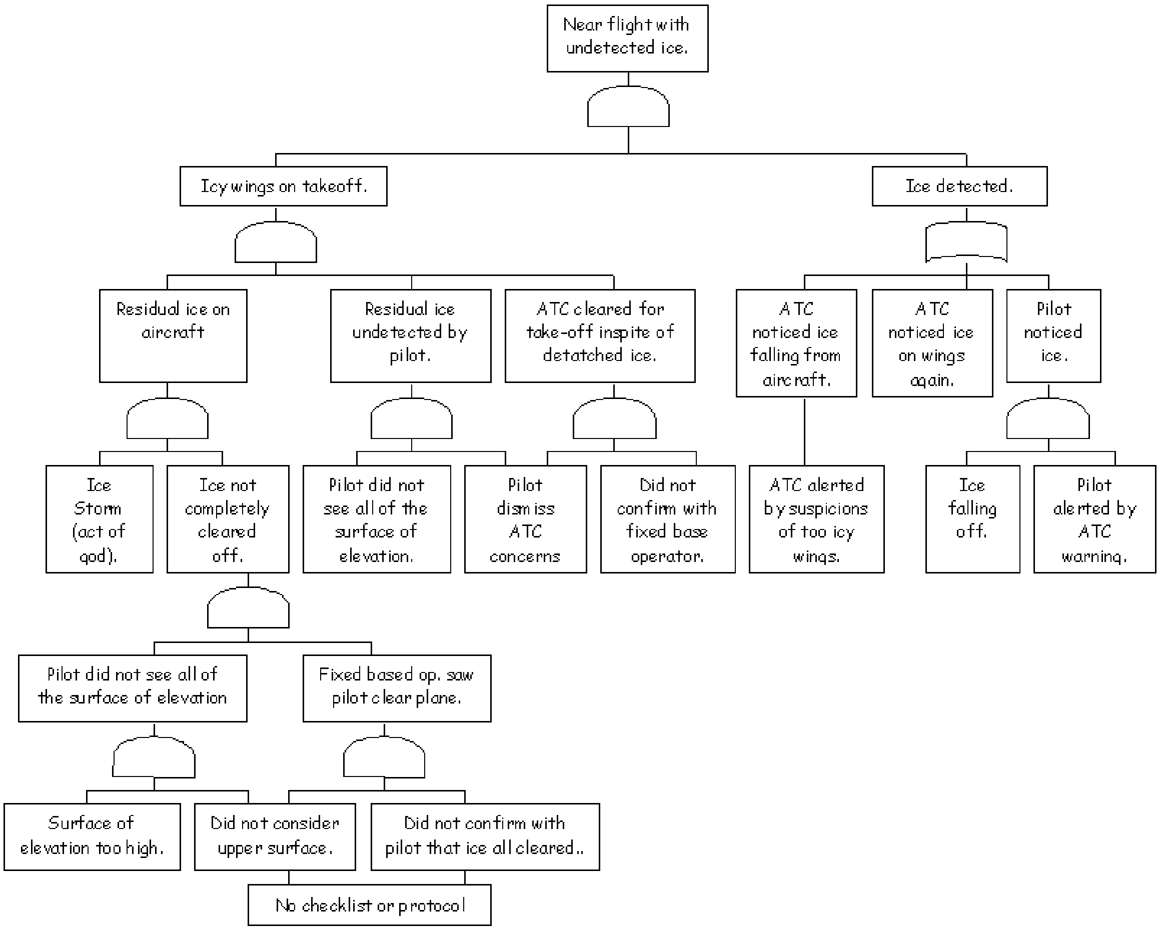
Near flight with undetected ice.
Icy wings on takeoff.
Ice detected.
Residual ice on aircraft
Residual ice undetected by pilot.
ATC cleared for take-off inspite of detatched ice.
ATC noticed ice falling from aircraft.
ATC noticed ice on wings again.
Pilot noticed ice.
Ice Storm (act of god).
Ice not completely cleared off.
Pilot did not see all of the surface of elevation.
Pilot dismiss ATC concerns
Did not confirm with fixed base operator.
ATC alerted by suspicions of too icy wings.
Ice falling off.
Pilot alerted by ATC warning.
Pilot did not see all of the surface of elevation
Fixed based op. saw pilot clear plane.
Surface of elevation too high.
Did not consider upper surface.
Did not confirm with pilot that ice all cleared..
No checklist or protocol

Figure 11.20: Causal Tree from McElroy's Evaluation of PRIMA (2)

during by the second participant:

- Operating procedures. This is represented by the node labelled OP in the Eindhoven Classification Model of Figure 11.9. The participant identified that the incident was the result of inadequate procedures. This category was also identified by the first participant for the helicopter case study.

- System Status. This is shown as HK1 in the Eindhoven Classification Model. The participant identified that the operator did not have an accurate knowledge of the "state and dynamics" of the system at key points during the incident [529].

- Permit. This is represented by HR2 in the Eindhoven Classification Model. The participant identified that the operator failed to obtain a permit or licence for activities where extra risk was involved. This category was also identified by the first participant for the helicopter case study.

- Checks. This is represented by HR4 in the Eindhoven Classification Model. The participant indicated that the operator had failed to conduct sufficient checks on the local system state to ensure that it complies with the expected conditions.

- Planning. This is shown as HR5 in the Eindhoven Classification Model. The participant identified that the activity was not fully planned. Appropriate methods were not identified nor were they carried out in a well-defined manner. This category was also identified by the first participant for the helicopter case study.

- Unclassifiable behaviour. This is shown as X in the Eindhoven Classification Model. The participant also identified that some of the causal factors denoted in Figure 11.19 could not be classified using the model. This category was also identified by the first participant for the helicopter case study.

McElroy's initial analysis had also identified Checks (HR4), Planning (HR5) and Unclassified behaviour (X) as root causes for this incident. The other categories were omitted. In addition, McElroy also identified License (HR1) as a causal factor. He argued that the operator in question must be qualified to do the job. He also identified Management Priorities (MP) as an issue in this incident. Top or middle management placed pressure on the operator to deviate from recommended or accepted practice. As noted in previous sections, it is difficult to reconstruct the thought processes that either of the participants used to justify their categorisation. McElroy notes in several places that the participants lacked the additional information that would have supported hypotheses about, for instance, the organisational causes of an incident. These are intriguing results. McElroy's results perhaps reflect the participants' suspicions that there must have been organisational causes to explain the operators' behaviour. If this is true then perhaps we are experiencing the consequences of the recent emphasis on the managerial and organisational causes of failure. These will be diagnosed as potential causes even when investigators are not provided with sufficient evidence to confirm these potential causes!

A number of methodological criticisms can be raised about McElroy's study. As mentioned, the lack of alternative data sources often forced the participants to make inferences and assumptions about potential causal factors. This led to causal trees and root cause classification that resembled rough 'sketches' of an incident. There criticisms can be addressed by acknowledging the severe time constraints that affected McElroy's work. They can also be countered by arguing that these high-level interpretations may resemble the level of detail that can be derived from an initial analysis of an incident report prior to a secondary investigation. It also provides an accurate impression of the 'rough' forms of causal analysis that can be conducted for contributions to anonymous incident reporting systems. In such circumstances, investigators are also constrained by the information sources that are available to them without compromising the identity of the contributor.

Further objections can be raised about the lack of empirical support for McElroy's work. He does not attempt to quantify agreement between his own causal analysis or that of the other participants. Given the limited data that he was able to obtain, this should not be surprising. He does not, however, speculate on measures that might be used. These is a vast range of techniques that can be used to represent and compare the structure of arbitrary tree structures [450, 451]. These algorithms might be used to detect patterns within the structure of causal trees. For example, Lekberg argues that an investigator's education background can bias their causal analysis [484]. Similarity metrics, for example based on vector representations, might be used to determine whether investigators from similar educational backgrounds produce measurably similar tree structures for the same incidents.

There are certain ironies about McElroy's approach. He framed his study as an experimental comparison between his own analysis and that of participants who were trained in incident analysis. he controlled the materials that were available to the participants and gave them the same training in the PRISMA technique. Having established these conditions, he lacked the resources to perform the additional tests that would have thrown light on many important issues. For instance, he did not counter-balance the incidents that were presented to the participants. This makes it difficult to determine whether any observed effects stem from the participant or the incident being studied. Similarly, McElroy only obtained access to two trained participants. Such a sample is inadequate to support any general conclusions. It should be stressed, however, that McElroy views his study as an initial experiment. It was intended to act as a marker for future studies that might attempt to assess whether investigators can *use* a causal analysis technique rather than just assessing their subjective attitudes towards someone else's application of an approach, as Munson had done [552].

This section has summarised recent attempts to assess the strengths and weaknesses of different causal analysis techniques. We have seen that these studies have only provided preliminary results.

The main conclusion from all of the work that we have cited in this section is that further research is needed to validate the many benefits that are claimed for the approaches that have been summarised in this chapter. It is, however, also possible to identify a number of more specific conclusions that might guide the future validation of causal analysis techniques:

- *Consider a broad range of stakeholders.* Previous studies have almost exclusively focussed on the investigators' assessment of causal analysis techniques. This is natural given that they are likely to determine whether or not a particular approach can be applied in the field. It should not be forgotten, however, that there are many other groups and organisations that must participate in, or validate, incident investigations. For instance, Chapter 3 discussed the role that regulators play in many reporting systems. A technique that satisfies investigators but does not convince regulatory bodies is unlikely to be acceptable. Similarly, it is important that any potential causal analysis technique should also be acceptable to those who must pay for its application. If this is not the case then there will be continued and increasing pressure either to reject the approach or to 'cut corners' in order to save resources.

- *Consider longitudinal factors as well as short-term effects.* All of the studies that we have presented are based around relatively short-term assessments of particular techniques. In particular, Munson and McElroy's evaluations took place over several hours. They do not, therefore, provide much evidence about the long-term benefits that might be provided by the consistent application of causal analysis techniques. There are also a range of detailed issues that are difficult to examine without more sustained studies. For instance, it is often necessary for investigators to revisit an analysis at some time after it was originally conducted. They may want to determine whether or not a subsequent incident has the same causal factors. In such circumstances, it is important not simply to identify the results of a causal analysis. It is equally important to understand the reasons *why* a particular decision was reached.

- *Consider the range of incidents in an evaluation.* It can be difficult to ensure that any assessment presents its participants with an adequate range of incidents. If this is not done then the utility of a causal analysis technique may be demonstrated for a sample of incidents that do not reflect the problems that are reported to the sponsoring organisation. There are further aspects to this issue. It may not be sufficient to base an evaluation on an accurate sample of current incidents. Given the overheads associated with training staff and financing the implementation of a new causal analysis technique, it is likely that any approach will be used for a significant period of time. If this is the case then any validation must also consider whether incidents will change during the 'lifetime' of an analysis technique. For example, Chapter 3 has argued that the increasing integration of computer-controlled production systems is posing new challenges in forensic software engineering. None of the techniques presented here, with the possible exception of WBA, explicitly addresses these challenges [411].

- *Consider the impact of individual or team-based investigation.* The studies of Munson, Benner and McElroy focussed on the performance and subjective assessments of individual investigators. Munson even went out of his way to prevent 'collusion' between the participants in his study. In contrast, Van Vuuren's evaluation involved teams of engineers, domain specialists, managers and safety experts. This reflects his intention to assess the application of this technique without the usual experimental controls that were accepted by Munson and McElroy. It is difficult, however, to determine whether team composition had any effect on the causal analyses reported by Van Vuuren. His published work says remarkably little about these issues. Work in other areas of groupwork have indicated that such factors can have a profound impact upon the successful application of design and analysis techniques [489, 556]. For example, the ability to use drawings and sketches as a medium of negotiation and arbitration can have a powerful effect during group confrontations. Attention may be focussed more on the shared artifact and less of the individuals involved in the discussion. We do not know whether these effects are also apparent in the application of causal analysis techniques.

- *Consider causal analysis in relation to other phases of investigation.* Benner reiterates the importance of evaluating any analytical technique within the context of a wider investigation

[73]. Analysis techniques are unlikely to yield sufficient explanations if investigators have not been able to elicit necessary information about the course of an incident. This argument was first made in the context of accident investigations. Unfortunately, cost limitations and the constraints of confidentiality/anonymity can prevent investigators from obtaining all of the data that they may need to complete a causal analysis. All of the techniques introduced in this chapter, with the exception of PRISMA, were developed to support accident investigations. These are, in one sense, information rich environments. In contrast, the particular characteristics of incident reporting systems may make relevant information very difficult to obtain. Any assessment must not, therefore, provide participants with information that they might not otherwise have available during the application of a particular technique.

- *consider which stage of an investigation is being assessed.* As we have seen, McElroy's initial evaluation of the application of an analysis technique produced results that were compatible with the early stages of an investigation. The participants produced trees that 'sketched' the outline of an incident. They did not produce polished artifacts that might provide consistent and sufficient causal explanations. Techniques that are intended to provide such quality control must, therefore, be validated in a way that enables investigators to achieve some degree of competence in the more 'advanced' use of the approach.

This is not an exhaustive list. Previous attempts to validate particular approaches have done little more than to sign-post areas for further work. It is equally important not to underestimate the importance of the small number of pioneering studies that have begun to validate the claimed benefits of causal analysis techniques.

11.5 Summary

This section has reviewed a broad range of techniques that can be used to support the causal analysis of safety-critical incidents. The opening sections build on our application of ECF analysis in Chapter 10 by introducing alternative event-based techniques. The related approaches of Multilinear Event Sequencing (MES) and Sequentially Timed and Events Plotting (STEP) were presented. These techniques all encourage analysts to use semi-formal, graphical or tabular notations to construct causal models of the events that lead to particular incidents. These notations provides great flexibility; investigators have considerable freedom in the manner in which they construct a causal model. Counterfactual reasoning is then, typically, applied to identify root causes from the candidate causal factors that are represented in a semi-formal model. Unfortunately, the flexibility offered by these approaches can also be a weakness. There are few guarantees that different investigators will derive the same results using this approach. Similarly, it is also unlikely that the same investigator will be able to replicate the details of their analysis at a later date.

Event-based techniques were, therefore, contrasted with approaches that exploit check-lists. These techniques provide investigators with a restricted choice of causal factors. Management Oversight and Risk Tree (MORT), Prevention and Recovery Information System for Monitoring Analysis (PRISMA) and Tripod all exploit variants of this underlying idea. The enumeration of causal factors guides and prompts investigators. It can also help to encourage consistency in an analysis. This is particularly important if national or regional comparisons are to be made between the causal factors of incidents that occur at a local level. Aggregated statistics would be unreliable if different investigators identified different causal factors behind the same incident. Of course, the price of consistency is that it may be difficult to identify an appropriate causal factor from the list of choices that are offered by these techniques. MORT and PRISMA address this potential caveat by encouraging investigators to extend the basic enumerations to reflect regional or domain-dependent variations in the incidents that are reported.

A further limitation of checklist approaches is that it can be difficult to check whether a particular analysis provides a consistent or sufficient explanation of a safety-critical incident. This chapter, therefore, introduced a range of formal causal analysis techniques. These approaches exploit mathematical systems of reasoning and argument to provide clear and concise rules about what can and what cannot be concluded about the causes of an incident. In particular, we have introduced WBA, partition techniques for non-deterministic causation and Bayesian approaches to subjective, probabilistic causation. Although these techniques are not widely used, they offer a number of potential benefits. They avoid many of the limitations that others have identified for the existing techniques that we have introduced in previous paragraphs [453, 482]. The rules that govern the application of these techniques provide objective criteria for verifying that an analysis is correct. The importance of ensuring the consistency and completeness of any analysis is also increasing significant given the rising cost of litigation in the aftermath of adverse occurrences. The modular approach supported by WBA and partition methods provides one means of addressing the increasing complexity of many safety-critical incidents. These benefits will only be realised if we can develop techniques that will enable non-formalists to participate in their application. At present, it can be difficult for those without an extensive background in mathematics to understand the strengths and the limitations of a particular formal analysis. Fortunately, many of the underlying mathematical models that support these causal analysis techniques can also be incorporated into software tools. There is also considerable potential for developing graphical and tabular representations that can be used to communicate more formal aspects of a causal analysis.

This chapter went on to describe attempts to validate some of the causal analysis techniques that we have described. Van Vuuren conducted bottom-up studies that involved the implementation of the PRISMA approach within several different industries [844]. He was then able to perform a comparative analysis of the different role that organisational factors played in a variety of different contexts. He did not, however, perform a detailed analysis of investigators' experiences in applying the causal analysis technique. In contrast, Benner provided a generic set of criteria that can be applied in a top-down manner to assess different accident models and methodologies [73]. By extension these same criteria might also be applied to assess different approaches to causal analysis. He relied largely upon his own subjective assessments. Munson, therefore, recruited an expert panel of fire fighters to apply similar criteria to a case study that had been analysed using Fault Trees, STEP and Barrier Analysis [552]. He was able to replicate results that suggested there were strong subjective preferences for STEP and Fault Trees over Barrier Analysis. Unfortunately, this study did not demonstrate that potential investigators might be able to apply any of these techniques themselves. McElroy, therefore, combined elements of Van Vuuren and Munson's approach when he asked a panel to apply the causal trees and Eindhoven Classification Model of the PRISMA technique [529]. This study revealed striking differences between the manner in which some people have proposed that causal analysis techniques should be used and the way in which investigators might actually use them in the early stages of an investigation. Rather than a detailed and careful analysis of the causal factors leading to an incident, the participants used them to sketch high level causes. They were less concerned with the consistency and sufficiency of an argument than they were with providing a clear overview of the incident itself. This, in part, reflects the important point that causal analysis techniques may have to play a variety of different roles during different stages of an investigation.

Our analysis of previous attempts to validate causal analysis techniques has revealed how little we know about the comparative strengths and weaknesses of these different approaches. We know from recent reports that current techniques are failing to support investigators tasks in many industries [482, 453]. This is another area that requires considerable further applied research so that practitioners can have greater confidence in the methods that are being proposed. The importance of this point cannot be underestimated. Several research teams are currently developing 'systemic' approaches to the causal analysis of incidents and accidents. These techniques are intended to address the challenges that are being posed by the failure of increasingly complex, tightly coupled systems. Unfortunately, less attention has been paid to the problem of demonstrating the practical benefits of these techniques than is currently being invested in their development.

It is worth emphasising that increasing complexity is one of several challenges that must be addressed by novel analysis techniques. They must also be proof against the sources of bias that influence the findings of many reports. Ultimately, it is not enough to show that any analysis technique can 'correctly' identify the causes of an incident. It must also demonstrate that it cannot easily be used to identify 'incorrect' causes. This is a significant burden given the many different forms of bias that might affect a causal analysis:

1. *Author bias.* This arises when individuals are reluctant to accept the findings of any causal analysis that they have not themselves been involved in.

2. *Confidence bias.* This arises when individuals unwittingly place the greatest store in causal analyses that are performed by individuals who express the greatest confidence in the results of their techniques. Previous work into eye-witness testimonies and expert judgements has shown that it may be better to place greatest trust in those who do not exhibit this form of over-confidence [223, 759].

3. *Hindesight bias.* This form of bias arises when investigators criticise individuals and groups on the basis of information that may not have been available to those these participants at the time of an incident. More generally it can be seen as the tendecy to search for human error rather than deeper, organisational causes in the aftermath of a failure.

4. *Judgement bias.* This form of bias arises when investigators perceive the need to reach a decision within a constrained time period. The quality of the causal analysis is less important that the need to make a decision and act upon it.

5. *Political bias.* This arises when a judgement or hypothesis from a high status member commands influence because other respect that status rather than the value of the judgement itself. This can be paraphrased as ‘pressure from above’.

6. *Sponsor bias* . This form of bias arises when a causal analysis indirectly affects the prosperity or reputation of the organisation that an investigator manages or is responsible for. This can be paraphrased as ‘pressure from below’.

7. *Professional bias* . This arises when an investigators’ colleagues favour particular outcomes from a causal analysis. The investigator may find themselves excluded from professional society if the causal analysis does not sustain particular professional practices. This can be paraphrased as ‘pressure from beside’.

8. *Recognition bias.* This form of bias arises when investigators have a limited vocabulary of causal factors. They actively attempt to make any incident ‘fit’ with one of those factors irrespective of the complexity of the circumstances that characterise the incident.

9. *Confirmation bias.* This arises when investigators attempt to interpret any causal analysis as supporting particular hypotheses that exist before the analysis is completed. in other words, the analysis is simply conducted to confirm their initial ideas.

10. *Frequency bias.* This form of bias occurs when investigators become familiar with certain causal factors because they are observed most often. Any subsequent incident is, therefore, likely to be classified according to one of these common categories irrespective of whether an incident is actually caused by those factors [394].

11. *Recency bias.* This form of bias occurs when the causal analysis of an incident is heavily influenced by the analysis of previous incidents.

12. *Weapon bias.* This form of bias occurs when causal analyses focus on issues that have a particular ‘sensational’ appeal. For example, investigators may be biased to either include or exclude factors that have previously been the focus of press speculation. Alternatively, they may become fixated on the primary causes of an incident rather than secondary causes that may determine the severity of an incident. For example, an investigation may focus on the driver behaviour that led to a collision rather than the failure of a safety-belt to prevent injury to the driver. This is a variant on the weapon focus that is described by studied into eye-witness observations of crime scenes [758].

The elements of this list illustrate the point that the success or failure of a causal analysis technique is, typically, determined by the context in which it is applied. For example, investigators can (ab)use causal analysis techniques by constructing causal chains that support particular, pre-determined conclusions. Such practices can only be discouraged by peer review during the various stages of a particular technique and by offering investigators a degree of protection against the sources of bias listed above. It should also be emphasised that causal analysis techniques are only one component in an incident reporting system. We cannot, therefore, assess the success or failure of such a system simply in terms of the sufficiency and completeness of the causal analyses that it produces. Such a validation must consider the success or failure of the recommendations that are justified by any causal analysis. These issues are addressed in the next chapter.

Chapter 12

Recommendations

Chapter 6 described how operators must often make immediate recommendations in the aftermath of an incident. These are intended to preserve the short-term safety of application processes. These immediate actions often exacerbate the consequences that they are intended to mitigate. numerous potential problems can prevent an effective response to an incident. Inadequate training, poor situation awareness, time pressure, the lack of necessary information, inadequate system support, pressure to preserve levels of service all impair operators' attempts to rectify an adverse situation. Chapter 6 also described a number of incident and emergency management techniques that are intended to reduce the impact of these factors. This chapter looks beyond the short-term recommendations that are made in the aftermath of an incident. In contrast, the intention is to examine the range of techniques that have been developed to identify potential remedies for the various causes that can be extracted from the approaches that have been introduced in Chapters 10 and 11.

12.1 From Causal Findings to Recommendations

A number of problems make it difficult to identify recommendations that reduce the likelihood or mitigate the consequences of future failure. The following paragraphs briefly summarises these problems. For example, there is a danger that investigators will continue to rely upon previous recommendations even though they have not proved to be effective in the past. Many authors have identified a form of 'conservatism' that affects large and complex organisations. It can take significant periods of time for new solutions to be adopted even when there is a considerable body of evidence that indicates the efficacy of alternative remedies.

There are several variations on the previous requirement. Many accidents occur because previous incidents have resulted in recommendations that do not adequately address the causes of previous incidents. Previous incidents provoke a range of different recommendations that reduce particular types of failure but which do not target underlying safety problems. Often these piecemeal recommendations avoid the expense or political involvement that are eventually committed in response to a subsequent accident. This can be illustrated by the US Central Command's investigation into a 'friendly fire' accident at Udairi Range in Kuwait. Previous incidents had resulted in procedures that were intended to ensure that crews were prevented from deploying their weapons if there was any danger of them mistaking observation posts for potential targets. Four previous incidents involving similar close-support operations had led to a number of local remedies being taken to minimise any potential confusion. Range procedures for fixed-wing aircraft to ground operations were changed to restrict delivery of ordnance to within two kilometers of Bedouin camps. The target was also altered to decrease the chances of any confusion. A tower was also constructed to help distinguish an observation post. The rooftop of the tower was painted white with a red cross. All of these physical changes failed, however, to address the overall problems of ensuring that crew did not inadvertently deploy their weapons at an observation post. The report concluded that 'despite four documented incidents in the past eight months, and attempts to improve conditions, observation posts and targets remain hard for pilots to see day or night' [824].

It is important that investigators should avoid arbitrary or inconsistent recommendations. Similar causal factors should be addressed by similar remedies. Of course, this creates tensions with the previous guidelines. The introduction of innovative solutions inevitably creates inconsistencies. The key point is, however, that there should be some justification for treating apparently similar incidents in a different manner. These justifications should be documented together with any recommendations so that other investigators, line managers and regulators can follow the reasons why a particular remedy was proposed.

Different organisations have proposed radically different approaches to the influence of financial or budgetary constraints on the identification of particular recommendations. Some organisations, such as the US Army have argued that 'the board should not allow the recommendation to be overly influenced by existing budgetary, material, or personnel restrictions' [806]. Other incident reporting systems, such as the local hospital systems that have been mentioned in previous chapters, accept a more limited horizon in which any recommendations must 'target the doable' [119]. The key point here is that investigators must ensure that their recommendations are consistent with the scope of the system. In the Army system, incident reporting is more open-ended with the implicit acknowledgement that significant resources may be allocated if investment is warranted by a particular incident. In the local system, incident reporting is constrained to maximise the finite resources of the volunteer staff who have run these systems. It is easy to criticise this constrained approach by recommending a more ambitious scope for the recommendations of a reporting system. It should be noted, however, that these systems have continued to introduce safety innovations for over a decade and without the national resources that are now being devoted to clinical incident reporting.

It is clearly important that any potential remedies must not introduce the possibility of new forms of failure into a system. Of course, it is easier to state such a requirement than to achieve it. The implementation of particular recommendations can introduce new forms of working that may have subtle effects. Given the relatively low frequency of many adverse occurrences it may only be possible to witness the safety-related consequences of those effects many months after particular recommendations have been introduced. The debate surrounding the concept of risk homeostasis provides numerous examples of such recommendations. Users may offset the perceived safety benefits of new regulations and devices against particular performance improvements. Cyclists who are compelled to wear safety-helmets may cycle faster than those who are not. Motorists who are provided with advanced braking systems may delay deceleration maneuvers [371, 370]. Others have conducted studies that reject the existence or the magnitude of such effects [532, 865]. The controversial nature of such studies not only indicates the difficulty of validating such effects, it also indicates the difficulty of ensuring that particular recommendations do not have any undesirable side-effects.

As mentioned above, the relatively low frequency of many adverse occurrences makes it difficult to determine whether or not recommendations have any palliative effect upon the safety of a complex application. In consequence, the impact of many recommendations must be measured through indirect means. For instance, it can be difficult to determine whether training operators in the potential causes and consequences of poor situation awareness can reduce the number of incidents that stem from this particular human factors problem. Simulator studies can, however, be used under restricted experimental conditions to show the short-term benefits of this training [864]. If such results cannot be obtained then there is a danger that the justification for any recommendation will be challenged. The support for particular remedies can be eroded as the salience of an incident fades over time. Further support for a recommendation can be elicited by repeating measurements that demonstrate the benefits of a particular approach. Unfortunately, these indirect measures can also be used to justify particular recommendations even when there is more direct evidence that casts doubt on the usefulness of a particular approach [410].

Previous sections have argued that incidents seldom recur in exactly the same manner. Future failures are often characterised by subtle variations in the causal factors that have been identified in previous failures. It is, therefore, important that recommendations are proof against these small differences. Similarly, recommendations must be applicable within a range of local working environments. They must protect against similar failure modes even though individual facilities may exploit different technical systems and working practices. These problems are, typically, addressed by drafting guidelines at a relatively high-level of abstraction. Safety managers must then interpret these recommendations in order to identify the particular remedies that might prevent future failures. If guidelines are too context specific then it can be difficult for safety managers to identify those lessons that might be usefully transferred to their own working environment.

There is, of course, a tension between identifying recommendations that protect a diverse range of systems and drafting recommendations that provide safety managers with the level of detail that is necessary to guide subsequent intervention. If recommendations are drafted at a high level of abstraction then there is a danger that different managers will choose to interpret those recommendations in ways that reflect arbitrary preferences rather than the local operating conditions, mentioned above. This is illustrated by a recent US Army safety alert. Previous incidents had led to ground precautionary messages that recommended military personnel not to use commercial heaters in unventilated areas 'use of unflued or unvented heaters is inherently dangerous because they vent exhaust containing carbon monoxide into living spaces' [815]. However, many soldiers chose to disregard this generic warning and continued to use heaters inside tents. These were not well ventilated The fabric was not intended to 'breathe' and several soldiers died as a result.

Recommendations are, typically, made by an investigator to the statutory body that commissions their work. These statutory organisations must then either accepting the recommendations or explain the reasons why they choose to reject them. If a recommendation is accepted then the statutory or regulatory body must ensure that it is implemented. This division of responsibilities is apparent in the aviation [423], maritime [833] and nuclear industries [204]. There are, however, occasions when investigators must identify who will be expected to satisfy a requirement. For instance, the US army requires that 'each recommendation will be directed at the level of command / leadership having proponency for and is best capable of implementing the actions contained in the recommendation' [806]. Such recommendations encapsulate good practice. If investigators do not identify suitable proponents for a recommendation then there is a danger that it may eventually be passed back and forth between a number of different organisations [444].

These potential difficulties make it important that any recommendations are well supported by the products of causal analysis. If this is not the case then it can be difficult for investigators to justify why particular remedies were, or were not, advocated in the aftermath of an incident. The following section, therefore, identifies a number of requirements that are intended to ensure that the results of a causal analysis can be used to support subsequent interventions.

A number of factors complicate the task of extracting appropriate recommendations from the findings of a causal analysis. For example, it can be difficult for investigators to assess the relative priorities of particular causal factors so that resources can be targeted towards effective forms of intervention. This task is further complicated because regional factors can reduce the impact of particular recommendations or, in extreme cases, can even mitigate any beneficial effects. Similarly, it can be difficult to identify recommendations that offer long term benefits rather than immediate or short-term palliatives. Finally, all of these problems are compounded by the difficulty of ensuring agreement amongst the diverse and multi-disciplinary groups that must concur with the recommendations that are produced by an investigation.

12.1.1 Requirements for Causal Findings

Many organisations deliberately 'target the doable' by restricting the focus of their recommendations to changes that affect the teams which support and implement a reporting system. In general, however, recommendations affect many different groups within complex organisations. One consequence of this is that representatives of these diverse interests must participate in the identification of remedial actions. At the very least, they must consent to their implementation. This can create a number of pragmatic concerns. For instance, recommendations may be drafted to address causal findings that were identified using one of the analysis techniques introduced in previous chapters. The products of some of these techniques, including Why-Because Analysis (WBA) and non-deterministic models of causation, cannot easily be understood by non-mathematicians. It is, therefore, important that the findings of any causal analysis are translated into a form that is readily accessible to those who must participate in and consent to the identification of recommendations in the aftermath of an incident. The following paragraphs summarise a number of further requirements that help in the use of causal analysis techniques to guide remedial actions.

Summarise the nature of the incident

In order to understand the significance of the causal findings that guide particular recommendations, investigators must summarise the course of a safety-critical incident. The US Army's Accident Investigation and Reporting Procedures Handbook requires an explanation of 'when and where the error, material failure, or environmental factor occurred in the context of the accident sequence of events; e.g., during preflight, while driving, etc' [806]. Information is also required to identify the individuals who are involved in an incident by their duty position or name. Components must be unambiguously denoted by a part or national stock number. Any contributory environmental factors must also be described. These requirements are often codified in the fields of an incident reporting forms. Operators are required to state the 'national stock number' of a failed component. They may also be asked to provide information about potential environmental factors and so on. As we have seen, however, the information that is provided by an initial report can also be supplemented by subsequent reconstructions. Chapters 8 and 9 introduced a number of techniques, including computer-based simulation, that can be used to model the course of an incident. These techniques underpin the causal analyses that were described in Chapters 10 and 11.

Summarise the causal findings

It is important that the products of any causal analysis are accessible to those without any formal training in the techniques that were used to identify them in the first place. This may seem like an unrealistic requirement given the underlying and inherent complexity of some causal models. It is, however, a prerequisite for ensuring a broad participation in the identification and implementation of any subsequent recommendations. If this requirement is not satisfied then there is a danger that other investigators, safety managers, regulators or line managers will mis-interpret the findings of any STEP analysis, PRISMA categorisation or Tripod modelling. It is for this reason that WBA employs a graphical form that can include natural language annotations in addition to the clausal forms of the more formal analysis. The US Army summarises the requirement to provide the following details:

> "For human error, identify the task or function the individual was performing and an explanation of how it was performed improperly. The error could be one of commission or omission; e.g. individual performed the wrong task or individual incorrectly performed the right task. In the case of material failure, identify the mode of failure; e.g. corroded, burst, twisted, decayed, etc. Identification of the directive (i.e. Maintenance / technical manual, SOP, etc.) or common practice governing the performance of the task or function. e. An explanation of the consequences of the error, material failure, or environmental effect. An error may directly result in damage to equipment or injury to personnel, or it may indirectly lead to the same end result. A material failure may have an immediate effect on equipment or its performance, or it may create circumstances that cause errors resulting in further damage / injury inevitable. Identification of the reasons (failed control mechanisms) the human, material, environmental conditions caused or contributed to the incident. A brief explanation of how each reason contributed to the error, material failure, or environmental factor." [806]

This quotation is interesting for a number of reasons. In particular, it provides an abbreviated checklist for the causal factors that must be considered when analysing particular types of failure. For instance, any analysis must consider the particular mode that characterised a materials failure. Such high-level guidance provides a lightweight means of combining the benefits of checklist approaches, such as MORT, with the more open form of causal analysis, encouraged by STEP and MES. The previous quotation urges investigators to consider the control mechanisms that caused or contributed to an incident. This is interesting because it acts as a reminder to consider critical aspects of an analysis even if investigators choose not to exploit barrier analysis or the related concepts in Tripod.

Explain the significance of causal findings

Chapter 7 introduced contextual details, contributory factors and root causes. Subsequent chapters have described a range of further distinctions that have been introduced by both researchers and by practitioners. These include concepts such as proximal and distal causes [701], particular and general causes [508], deterministic and stochastic causes [313] and so on. Irrespective of the precise causal model that is adopted, it is important that investigators provide some indication of the perceived importance of any particular causal finding. For instance, the US Army recommends that findings are categorised as 'Found; Primary Cause, Found; Contributing, Found; Increasing Severity of Damage/Injuries, or Found; Not Contributing' [806]. As before, this recommendation acts as an important reminder to incident investigators. For example, the previous chapter briefly summarised the potential impact of weapon bias. Investigators can become fixated on the primary cause of an incident at the expense of secondary failures that increased the severity of any outcome. By explicitly reminding investigators to consider these factors, these guidelines encourage analysts to look beyond the driver behaviour that leads to a collision. They are, for instance, encouraged to identify the reasons why a safety-belt failed or why the emergency response was delayed. The same guidelines also encourage investigators to separate the presentation of primary causes from contributory factors by noting that 'THE FINDING(S) LISTED BELOW DID NOT DIRECTLY CONTRIBUTE TO THE CAUSAL FACTORS INVOLVED IN THIS INCIDENT; HOWEVER, IT (THEY) DID CONTRIBUTE TO THE (SEVERITY OF INJURIES) OR (INCIDENT DAMAGES)' [806]. This quotation shows how it is important not only to consider the information that must be identified by any causal analysis but also the format in which that information is transmitted. The Army handbook requires that such 'contributing' causes can easily be distinguished form the 'primary' causes that directly led to an incident.

Justify excluded factors

Not only is it important to explain the significance of those causal factors that did contribute to an incident, it is also necessary to explain why particular 'causes' did NOT contribute to an adverse occurrence . Investigators must not only explain why recommendations address particular aspects of a system, they must also explain why those recommendation did NOT address other aspects of the system. These excluded causes fall into two categories. Firstly, those factors that did not cause or exacerbate this incident but which have the potential to cause future failures if uncorrected. As before, the products of this form of causal analysis must be clearly distinguished from 'primary' and 'contributory' causes: 'the findings and recommendations fitting this category will be separated from those that caused the incident or those that did not cause the incident but contributed to the severity of injuries / damage' [806]. There is, however, a second class of excluded 'causal' factors that must also be considered in the findings of any causal analysis. These are the factors that might have caused to, or exacerbated, an incident but which were considered not to be relevant to this or future failures. Without such justifications it is impossible for other investigators, for managers and for regulators to distinguish between such those factors that were considered but rejected and those that were never even considered in the first place.

Summarise the evidence that supports or weakens each finding

This book has repeatedly argued that investigators and analysts must justify and document the reasons why particular decisions are taken at each stage of their work. Without this additional information it can be difficult for other investigators, for regulators and for other statutory bodies to understand why an investigation proceeded in a particular manner. It can be difficult to follow the reasons why a secondary investigation was not initiated. It may be difficult to identify the factors that led investigators to commit resources for computer-based simulations in one incident and not another. Similarly, it can be hard to understand why resources were not allocated to support a detailed causal analysis. The US Army handbook recognises the need to justify the outcome of a causal analysis; "Each cause-related finding must be substantiated." [806] The cursory nature of this requirement is, perhaps, indicative of a wider failure to recognise the importance of such justifications. All too often, individuals and groups must endeavour to 're-live' their decision making processes during the course of subsequent litigation. A number of techniques can be used to document the justifications that support particular causal arguments. For instance, Chapter 9 introduced Conclusions, Analysis, Evidence (CAE) diagrams. These provide a means of linking the evidence that can be obtained in the aftermath of an incident to the arguments for and against a conclusion. These graphical structures are intended to provide a high-level overview of the justifications that support particular causal findings.

12.1.2 Scoping Recommendations

Causal findings help to guide the drafting of appropriate recommendations. The US Army handbook, cited in previous paragraphs, illustrates this relationship by advising that each finding is printed next to the remedy that has 'the best potential' for avoiding or mitigating the consequences of future incidents [806]. As we have seen, however, it can be difficult to identify appropriate recommendations. In particular, interventions must be pitched at the correct level. They must be detailed enough so that they avoid ambiguity. They must present the organisational, human factors and systems details that are necessary if future incidents are to be avoided. They must not, however, be so specific that the fail to capture similar incidents that share some but not all of the causes of previous incidents. The following paragraphs briefly describe some of the more detailed issues that must be considered when attempting to identify an appropriate scope for the recommendations in an incident report.

By time...

Previous sections have identified important differences between the short-term recommendations that are made in the immediate aftermath of an incident and the longer-term remedies that are, typically, the outcome of more considered investigations. Immediate instructions to alter operating practices may be supplemented by regulatory intervention to ensure that those changes are backed by appropriate sanctions. It is important to recognise, however, that very few recommendations ever provide indefinite 'protection' against future failures. In military systems this is best illustrated by the continuing problem of 'friendly fire' incidents.

	World War II 1942-1945	Korea 1950-1953	Vietnam 1965-1972	Desert Storm/Shield 1990-1991
Accidents	56%	44 %	54%	75%
Friendly Fire	1%	1%	1%	5%
Enemy action	43%	55%	45%	20%

Table 12.1: Battle and Non-battle casualties in the US Army [798].

Table 12.1 presents US Army figures for the changing impact of friendly fire incidents on army casualties in major combat operations since 1942 [798]. Such incidents, however, have a far longer history. One of the most (in)famous incidents occurred in April 1863 when Robert E. Lee's Army of Northern Virginia attempted to halt the Union Army of the Potomac's advance across the Rappahannock River near Chancellorsville. Lee left a small force to contain Major General Joseph Hooker and sent the remainder of his strength with 'Stonewall' Jackson to attack the Union flank. Jackson achieved considerable success and pushed ahead with a scouting party. As the party returned, they were mistaken for Union cavalry. Jackson was wounded and died soon after from complications that followed the amputation of his left arm [23]. Such incidents stem from a lack of situation awareness, often involving scouting parties and other advanced units. They also stem from the development of weapons that are effective at a range which is greater than the range at which combatants can easily distinguish friend or foe.

Such incidents were often seen as the result of undue recklessness, or bravery, on the part of the individuals involved. In the years following the Civil War, greater emphasis was placed on the development of communications systems and protocols that were intended to improve combatants' understanding of their combat situation. For instance, rules of engagement were drafted to identify situations in which it was 'safe' to engage a potential enemy. An example of such procedures can be found in the Rules of Engagement-Southeast Asia (U), JCSM-118-65, 19 February 1965 (Declassified 21 June 1988, NARA) which removed the US military's restriction against pursuit of Vietcong into Communist China. These required that hostile vessels could only be attacked in Vietnamese (RNV) or Thai territorial waters if it had been 'attacking or acting in a manner which indicates with reasonable certainty an intent to attack U.S./friendly forces or installations, including the unauthorised landing of troops or material on friendly territory' or 'engaged in direct support of attacks against RVN or Thailand'. Unfortunately, these new tactics and tools were not always successful in eliminating the problem of friendly fire. For example, the US Naval Institute published an account of an engagement during the Vietnam war. A B-57 from the 8th Bombardment Squadron attacked a US patrol boat after it had dropped its bombs on watercraft just north of demarkation zone. Coordination between the 7th Air Force and the naval forces was particularly poor. The Commander-in-Chief, Pacific, later observed that 'this incident is an apparent lack of tactical coordination between operational commanders'. The 7th Air Force investigation concluded that the vessel did not know the 'correct MAROPS challenge/response for air to surface'. The patrol boat had 'two means of identifying themselves to aircraft' using their running lights or by radio communications but 'the vessel did neither' [380].

'Friendly fire' accounted for some five percent of American casualties during Operation Desert Storm in 1991 [798]. These often had similar causes to incidents in previous conflicts. Many stemmed from communications problems. Deployment information was not passed along the chain of command. Other incidents again revealed the disparity between the range and effectiveness of modern weapons systems when compared to battlefield communications equipment. Following the gulf war, several initiatives started amongst allied armies to lessen the number of these incidents in future conflicts [747]. As can be seen, some of these initiatives focussed on new technologies. Others, however, have more direct parallels with the techniques that were proposed in previous conflicts:

- "Systems that align with the weapon or weapon sight and are pointed at the intended target. The system 'interrogates' the target – a reply identifies it as friendly, otherwise it is identified as unknown.
- 'Don't shoot me' systems use the Global Positioning System and other similar data sources. An interrogation is sent in all directions containing the targeted position. Friendlies present in that position return a 'don't shoot me' response.
- Situational awareness systems rely on periodic updates of position data to help users locate friendly forces.
- Non-cooperative target recognition systems compute a signature using acoustic and thermal signals, radio emissions, and other possible data sources. The system compares the signature in its library database to characterise the target as potentially a friend, foe or neutral." [22]

A number of reasons explain the way in which the similar hazards recur over time even though recommendations provide some immediate protection from a particular failure. For example, previous sections have cited research into risk homeostasis that determines whether or not users will sacrifice safety improvements in order to achieve other objectives. Car drivers will rely on advanced braking systems to save them from hazardous situations. A number of other potential problems can prevent previous recommendations from continuing to protect application processes. For instance, operators and managers may forget the importance of previous remedies as incidents and accidents fade from the 'group memory' [633, 635]. This process of 'forgetting' should not be underestimated given the relatively low frequency of many adverse occurrences. Organisational factors also intervene to increase the speed at which previous recommendations can be lost to those whose actions must be guided by them. The recommendations from less severe incidents may be lost more quickly than those of the ‘friendly fire’ examples, cited previously. The Canadian armed forces have one of the most advanced health and safety infrastructures of any military organisation. Their computer-based General Accident Information System automates the submission and partial filtering of incident and accident reports [148]. These reports are summarised in Safety Digests that are similar to the Aviation Safety Reporting Systems DirectLine publication, introduced in Chapter 5. They provide key personnel with ‘first-hand’ accounts of previous incidents. They also communicate the recommendations from investigations in a relatively accessible manner. Such feedback does not always have the strong, long-term remedial effect that might be expected. For instance, there have been several incidents involving the transfer of fuel under pressure between various types of bowser [137]. These have led military safety managers to stress that it is the “duty” of military personnel and civilian sub-contractors to refuse to engage in operations that jeopardise safety during peacetime [135]. Unfortunately, there recommendations have not had the impact that might have been hoped:

> “Training DND military/civilian personnel performing the transfer operation and those in the chain of command didn’t have experience or training to safely conduct this non-standard fuel transfer. Basic/advanced fuel handling training for National Defence military/civilian personnel requires further in-depth evaluation. In the interim, 19 Wing is conducting enhanced local training. The applicable engineering publications governing the safe handling of fuels are outdated, not accessible to all personnel and appear to be technically inferior to industry standards and other Air Forces’ publications.” [140].

This quotation provides several detailed reasons why many of the recommendations from incident reports are limited to a relatively short ‘shelf life’. Personnel may not have been provided with access to the initial information. They may have joined an organisation or have been re-deployed within an organisation well after the findings from an incident have been published. Staff may also be employed by sub-contractors who were not informed of the recommendations that were identified from previous incidents. Conversely, the organisation itself may have fallen behind best-practice in an industry. The previous quotation identifies that military procedures failed to meet civilian standards.

The long-term effectiveness of particular recommendations can also be undermined by changes in working practices. These need not reflect deliberate neglect or the failure to communicate the importance of adopting particular remedies in the aftermath of previous incidents. In contrast, these changes can be forced upon personnel by the introduction to new technologies. Some recommendations that ensure safe fuel transfer from bowsers can also be applied to other fuel storage mechanisms, such as bladders [135]. For instance, it is important to ensure that hoses are hydrostatically tested in both situations. Other recommendations cannot be directly transferred in this way. For example, previous bowser fires have established the importance of using industry-approved flow rates for particular fuel types [140]. This recommendation has some relevance for bladder devices. However, the particular properties of bladder devices require that recommendations from previous bowser fires must be carefully reinterpreted if they are to protect operators using these containers. Personnel must ensure that fuel is pumped to the bladder’s pressure rating rather than at its maximum filling speed.

Changes in the operating environment undermine previous remedies. For instance, many military organisations responded to ‘friendly fire’ incidents by implementing protocols, such as terms of engagement, that guide personnel on the actions to be taken before engaging a potential target. Battlefield communications systems have also been developed to help distinguish friend from foe. These remedies are tailored to meet the specific requirements of particular military organisations. It can be difficult to extend the same techniques to support joint operations by allied forces. For instance, there may not be the political support that is necessary to agree upon common terms of engagement. It is also rare for allied forces to share the same core communications technologies. One consequence of this is that joint operations often result in a large number of friendly fire incidents. Remedies that reduce incidents in particular scenarios may, therefore, not provide protection under changed operational circumstances.

Changing working practices, changing operational contexts and changing technologies create considerable problems for investigators who must ensure that their recommendations continue to protect the safety of a system and its operators. Several techniques have been proposed to reduce the impact of such changes on the remedies that are advocated by incident reporting systems. For example, investigators can explicitly specify the shelf-life of a recommendation. Any remedial actions need only be implemented until an end-date that is specified in the incident report. The regulatory or statutory body is then responsible for explicitly renewing any recommendation that might be made after the initial period of enforcement has expired. This approach has obvious disadvantages for any regulatory body that must constantly review a mass of relatively low priority recommendations. An alternative approach is to require that organisations periodically update their safety cases to ensure that they conform to recommendations that have been made since their previous appraisal. This review also provides an opportunity for companies to argue that previous recommendations may no longer hold given new working conditions or technological innovations. This approach also suffers from a number of limitations. For example, it can be difficult to identify an appropriate renewal period. Alternatively, companies may be required to revise a safety-case whenever new working practices or environmental conditions are introduced. Further technical difficulties complicate the task of updating a safety case, see for example [434].

Some incidents raise a variety of more 'pathological' temporal issues that exacerbate the problems of drafting and implementing appropriate recommendations. For example, the Singaporean army has made a number of recommendations that have reduced the number of heat related injuries reported in recent years:

> "During the first two days of heat exposure, light activities would be appropriate. By the third day of heat exposure, 3 kilometer runs at the pace of the slowest participant are feasible. Significant acclimatisation can be attained in 4 to 5 days. Full heat acclimatisation takes 7 to 14 days with carefully supervised exercise for 2 to 3 hours daily in the heat. The intensity of exercise should be gradually increased each day, working up to an appropriate physical training schedule adapted for the environment." [741]

As mentioned, these recommendations have encouraged a general decline in heat related injuries within the Singaporean defence forces since 1987. If we follow the argument that has been presented in previous paragraphs then it might be argued that greater concern should be devoted to other, potentially more pressing, safety recommendations. A number of factors have, however, combined to increase the salience of there recommendations. Since 1995 the army has continued to report approximately 3.5 cases per 1000 soldiers. These cases are not evenly distributed across all units. Training schools continue to suffer the highest incidence of heat-related injuried as new soldiers transition from civilian life. The political and social impact of these incidents is exacerbated by the Singaporean army's continued use of enlistment. In consequence, the 'shelf-life' or duration of a recommendation can be determined by a range of factors that may have relatively little to do with the relative frequency of particular incidents.

The previous example can be used to illustrate a number of further problems that complicate the task of drafting appropriate recommendations. For instance, the previous remedies are increasingly important at particular times in the year. The Singaporean army reports the highest number of heat injuries in April and May. This reflects increases in heat and humidity during those months. As we have seen, the salience of particular recommendations can decline when they are not perceived to be important to an operator's immediate task. In consequence, safety managers must make particular efforts to reinforce the importance of these guidelines during March and early April. The complexity of drafting appropriate recommendations is further complicated by the bimodal distribution of these incidents within the day. Peaks occur in the reporting patterns from 08:00 to 09:59 hrs and from 16:00 to 17:59 hrs. These peaks straddle the interval between 11:30 and 15:30 hrs during which formal physical training is prohibited according to Singaporean army regulations. Further complexity is introduced by the time-limits that determine appropriate mitigating actions. If an individual's heat exposure is less than 90 minutes then they should be offered plain, cool water during a recovery period. If the heat exposure exceeds 90 minutes then they should be offered a "cool, suitably flavoured carbohydrate-electrolyte beverage" with no more than 8%, or 2 table spoons of sugar per litre [741]. If the soldiers' heat exposure exceeds 240 minutes then they should be provided with a flavoured "carbohydrate-electrolyte beverage supplemented with one tea spoon of salt per litre".

By place...

The previous section has argued that it can be difficult to draft recommendations that can continue to have a medium or long-term effect on the safety of an application. Memories fade, working practices change and technology is seldom stable beyond the short-term. In consequence, regulatory or statutory intervention may be required to ensure continued compliance. Investigators may also be forced to draft their recommendations so that they are 'future proof' against these changes. Unfortunately, remedies that avoid reference to particular technologies and working practices are likely to be of little practical benefit. Lack of detail encourages ambiguity and safety managers may find it difficult to know how to implement remedial actions. These problems are compounded by the need to ensure that recommendations can be implemented in many different working environments that are often distributed across many different geographical locations. This is an increasing problem given recent initiatives to increase the coverage of national and international reporting systems. For instance, the initiatives of individual airlines led to the development of United States' Aviation Safety Reporting System (ASRS) in 1976. In the last five years this has, in turn, motivated attempts to establish a Global Aviation Information Network [308]. Similarly, medical reporting systems that were initially established in individual units within individual hospitals are now being extended to regional and national systems. For example, the UK's Royal College of Anaesthetists has introduced guidelines to encourage recommended practice in incident reporting within their specialism [715]. Both the UK Government [633] and Presidential initiatives [453] have advocated the expansion of these systems beyond the local and regional levels.

It is important to emphasise that national and international initiatives to expand the geographical coverage of incident reporting systems do not remove the need to draft recommendations that focus on particular local needs. For example, the Canadian Commander of the National Support Element Services Platoon and of Camp Black Bear in Bosnia-Herzegovina reported that the following actions had been taken to address previous safety recommendations:

> "First of all, we replaced three propane gas ranges in the kitchen and took steps to replace one tilting frying pan and procure another. The ranges in use were extremely old and beat up. In fact, the burners were cracked and the wire insulation was torn. New safety equipment in the kitchen at Camp Black Bear. Moreover, the plates inside the stoves had been removed, leaving the propane gas tubing unprotected. Hence, this equipment posed a serious risk to the people working in the food section. Our chief cook, Sgt Élément, is exceedingly proud of his new equipment." [143]

The importance of such local recommendations and actions cannot be exaggerated. They provide immediate and direct feedback to the individuals and groups who contribute to incident reporting systems. This is particularly significant for work groups that perceive themselves to be isolated from administrative centres. The Canadian units in Bosnia-Herzegovina provide a good example of groups who most need to be reassured that their potential problems are receiving prompt and direct attention.

There are also less obvious reasons for ensuring that recommendations address particular local concerns. If remedies are couched in abstract terms that can be applied to many different contexts they often lose the impact that can be observed from more direct and locally relevant recommendations. For instance, the previous actions might have addressed a requirement to 'review the safety of cooking appliances in all military camps'. Such generic recommendations can often be lost amongst the plethora of similar high-level guidance that is issued from 'lessons learned' systems. The introduction of particular local details can, arguably, provide more salient reminders even though the exact circumstances are not replicated in other working environments. This is an important feature of the anecdotes and 'war stories' that provide a critical learning resource for workers in a vast range of occupations. This analysis was confirmed during the interviews that help to form the EUROCONTROL guidelines for incident reporting in air traffic management [423]. Many controllers specifically asked that details about specific airports and shift patterns should be left in both the causal analysis and recommendations associated with individual incidents. They argued that this increased the perceived relevance of the analysis. These local details helped them to re-interpret particular recommendations within the often different context of their own working environment. The controllers who were interviewed continued to support these arguments even after it was suggested that such local details might compromise the anonymity of some reports. It remains to be seen whether the same opinions would be expressed by those individuals who are involved in an incident.

It is possible to identify a paradox that affects the drafting and implementation of recommendations from incident reporting systems. The effectiveness of such systems depends upon identifying remedies that can be applied well beyond the scope of the application or working group that first identified a potential problem. On the other had, drafting recommendations that can be applied beyond the context of particular working group often implies that investigators must strip out the contextual details that help operators understand the significance of an incident. There is also a danger that by expanding the scope of a recommendation, investigators will address propose remedies for problems that do not exist beyond the boundaries of a local system. This is a significant concern given that each recommendation is likely to carry significant costs in terms of the time and money that may be required to implement them. There is a further danger that by addressing these spurious recommendations, working groups may divert resources away from more critical remedies. Many organisations, therefore, impose triggering conditions that must be satisfied before an incident must be addressed by both local and regional recommendations. For example, the Canadian forces in the Former Republic of Yugoslavia investigated a total of 250 different topics during initial investigations into incidents during 1996. The 'Lessons Learned Information Warehouse' of the Army Lessons Learned Centre identified 142 possible areas of study but only 34 of these subjects were common to the majority of operations and rotations. There were identified using the following heuristics. Firstly, if an issue was identified during the early stages of the army's involvement in the region but ceased to be reported in latter stages then it was assumed that appropriate changes had been made and that no further recommendations need be made. This is a strong assumption. In other contexts, it would be necessary to seek further reassurance that such initial reports had been rectified. For example, the last heuristic represents such an approach. This is a further illustration of the temporal problems that arise when investigators attempts to assess the 'shelf-life' of a recommendation. Secondly, if an observation was made twice in the latter operations then it was retained as an issue for further investigation. The military reports do not state whether an issue would be retained if these reports were submitted by the same units on two different occasions. As with the previous heuristic, additional requirements might be necessary in other contexts to ensure that such generic issues did stem from more than one independent source. Finally, if an observation was made three times at any stage of the operations then it was retained as an issue. This increases the likelihood that high frequency reports, even in the early stages of the operation, will be addressed during subsequent investigations [134]

As mentioned, these criteria were used by the Canadian military as a means of filtering the 250 topics that were identified by individual reports from the units in the Former Republic of Yugoslavia. Only those issues that satisfied these various criteria were passed for the next level of analysis. In other words, they ceased to be considered as isolated examples that could be addressed by local recommendations. They were, in contrast, considered as more generic issues that required national or regional remedies. This distinction can be illustrated by the contrast between the general health and safety issues that affected Camp Black in the previous quotations and the following issues that emerged from this higher level analysis of more generic issues:

- "*Operations - Maps (Issue 26, page A-13):*
  Although units had adequate map coverage, the two map scales in use (1:50,000 and 1:100,000) did not coincide. Reporting of a grid on one scale produced an error in plotting on the other due to a difference in data.

- *Operations - Mines (Issue 32, page A-16):*
  Mines were the single largest producer of casualties on operations in the Former Republic of Yugoslavia. All [reports] indicated that units conducted extensive mine awareness training prior to and during the operation. Despite this training, the vast majority of mine incidents were directly attributable to a lack of situational awareness, understanding risks and recognising

> the indicators of a mined area." [134]

Many organisations operate similar filtering processes to those implemented by the Canadian defence forces. Incidents that are reported at a local level are collated and then assessed to determine whether they have regional or national significance. If an issue is considered to be sufficiently serious according to the filtering criteria, or subjective judgement of the gatekeeper [423], then more generic recommendations are drafted. This, typically, implies a process of abstraction that strips out the contextual details that have been noted in the previous citations. This can be illustrated by the US Army's response to repeated incidents involving poor situation awareness; 'many platoons continue to experience difficulty with situational awareness because they do not have a system in place to properly battle track and manage information' [800]. The perceived importance of this continuing problem ensures that any recommendation must be directed to all battlefield personnel rather than those who are engaged in regional operations. This contrasts with the previous quotation that focussed on incidents involving maps and mines that were specifically encountered by the Canadian element of the UN forces in the Former Republic of Yugoslavia. Incidents involving poor situation awareness cannot be viewed as local issues because similar problems led to the friendly fire incidents mentioned earlier. An analysis of previous incidents determined that battle tracking in platoon command posts failed to provide squad leaders with necessary details about enemy locations, friendly unit dispositions and the current state of combat operations in their area. Squad leaders, in turn, rarely provided sufficient detail for platoon leaders to gain a clear understanding of the significance and context of their objectives. The US Army responded to repated reports of similar incidents by directly considering situation awareness issues within its national training programmes:

1. "The platoons must provide brigade combat teams with the information necessary to have resolution of location, current status and missions of the Military Police units.

2. Military police platoons should be considered during the brigade combat team's clearance of fires drills. The platoon command post must track the current brigade operation to the resolution necessary to provide squad leaders with information to plan and conduct operations.

3. Prevent fratricide. The platoon command post must also disseminate and provide feedback on the Commander's Priority Intelligence Requirements and Critical Information Requirements. Platoon leaders must require squad leaders to submit timely situation reports and route reconnaissance reports." [800]

As can be seen, there recommendations have moved away from the specific problems encountered by particular units. They have also abstracted away from the more regional problems that are associated with a particular theatre of operations. In contrast, the recommendations are expressed in a generic manner that might be used to inform battlefield operations in any anticipated conflict. It is important to emphasise that investigators must be aware of the different strengths and weaknesses of recommendations that are pitched at a national rather than a local level. The benefits of extending the scope of any remedy are obvious. However, the sense of engagement that stems from addressing specific local concerns is difficult to obtain from this more generic approach. The previous recommendations could be aimed at any combat platoon in the US Army. The insights gained from the analysis of particular incidents are distilled into a format that resembles standard training manuals that lack the immediacy of more local approaches. This effect is more readily apparent in the recommendations that are intended to avoid administrative or financial 'incidents'. For instance, the following quotation is taken from the US Office of the Assistant Secretary of the Army's proposals to avoid problems in the acquisitions process:

> "People who show active hostility to changes are easy to spot and deal with. If they express their dissatisfaction honestly and openly, then their objections can be addressed. The agreement may be stronger for having resolved the points troubling such a person. Conflict resolution, after all, is one of the primary reasons for forming a partnering agreement. More difficult to deal with are those individuals who pay lip service to the partnering agreement while they quietly work against it. Their hostility is expressed with

> subtlety through stubbornness, procrastination, and inefficiency. While the agreement encourages actively working to find solutions to problems, a passive-aggressive person does nothing to further the process. Quite the contrary, they do whatever they can to wreck it. Reassign such a person to a job where they cannot block progress." [801]

Such advice is deliberately pitched at a very high level of abstraction. In consequence, it can appear to be little more than 'common sense'. Such advice can have potentially adverse consequences if offered to personnel who are faced with more direct and apparently pressing problems in their working lives [839].

The previous paragraphs have argued that there is a tendency for national and regional reporting systems to remove the local and contextual details that often increase the immediacy of particular incident reports. In consequence, recommendations can be seen as abstract requirements that have little relevance to more immediate problems. At worse, they can be resented as unwarranted impositions by external agencies that are intent on hindering the normal working practices of local teams. It is important to emphasise that this process of abstraction is not a necessary result of attempts to increase the scope of an incident reporting system. It is still possible to implement national and international systems that provide focussed information about detailed incidents. Unfortunately, this raises a number of fresh problems that must be addressed by investigatory and regulatory authorities. In particular, given that national and international systems may generate a large number of potential recommendations it can be difficult to ensure that particular members of staff can easily access all relevant recommendations. Several techniques have been developed to address this problem. Journals such as the ASRS' DirectLine or the Canadian National Defence forces' Safety Digest can publish information about individual incidents in 'key areas' that are selected by the staff who are responsible for running the system. These publications are then distributed to appropriate members of staff. Unfortunately, this approach can be extremely costly. Paper-based publications must, typically, be distributed to many different regions. It also relies upon investigators to identify a subset of incidents that should be publicised at a national level. The difficulty of this task increases in proportion to the scale of the reporting system.

Electronic publication techniques provide alternative means of providing key members of staff with access to the recommendations that affect their particular tasks. This avoids the problems associated with making an explicit decision only to publicise a small number of the insights that can be gained from a national reporting system. With appropriate tool support, this approach also avoids some of the overheads associated with the costs of updating and distributing paper based journals. This approach has been successfully exploited by a range of armed forces [801, 148]. Preliminary steps have also been taken to extend this approach as a means of encouraging international cooperation. The ABCA Coalition Operations Lessons Learned Database is a notable example of this approach. This database was established in 1999 as a joint venture between the American, British, Canadian and Australian armed forces. It was intended to "identify and resolve those key standardisation issues which would affect the ability of a military force, comprising two or more of the ABCA nations, to operate effectively and to the maximum ability of its combat power" [799]. The password-protected web-site provided user with the ability to perform full-text searches. They could also browse a full listing of documents by country of origin. Chapter 15 will describe some of the technological limitations that reduce the utility of this approach and will introduce a number of further solutions to these problems. For now it is sufficient to observe that there may be few guarantees that any particular member of staff will be able to access all of the recommendations that are relevant to their working tasks using computer-based systems.

It is important not to underestimate the problems that arise when attempting to draft recommendations that might usefully be applied across national boundaries. Previous chapters have argued that the increasing scope of an incident reporting system can result in a process of abstraction that hides contextual information. Cultural differences have the paradoxical effect of focusing international exchange almost exclusively on detailed technical issues. For example, it is difficult to translate previous advice on US Army acquisitions policies into cultures in which 'apparent acceptance and covert opposition' are acceptable and even anticipated forms of disagreement [878]. It is for these reasons that the exchange of safety-related recommendations can yield deep insights into the alliances that exist between national organisations. Cultural similarities arguably explain the

United Kingdom's participation in the ABCA coalition rather than a coalition with other European defence forces.

The previous analysis has argued that the effectiveness of an incident reporting system can be increased if investigators increase the scope of its recommendations. This, typically, involves abstracting away from local, contextual details so that lessons can be applied by operators working in different regions and even different countries. It is important to stress, however, that there are some situations in which there is a deliberate policy not to exchange information about safety related incidents. For example, the South African National Defence Force is still adjusting to the changes that were introduced when it was first made subject to both the Machinery and Occupational Safety Act, 1983 and the Occupational Health and Safety Act, 1994. For the first time, the Department of Defence has an explicit obligation to demonstrate compliance with the law, "or with the spirit of the law", in health and safety matters [708] Prior to the end of apartheid, there was a deliberate political motivation to promote self-sufficiency. This implied a willingness to learn from the mistakes of others but did not imply a willingness on the part of many governments to share those lessons. Even in the post-apartheid era, there are limits to the free exchange of information in military and strategic matters. This was implied in the recent White Paper on South African Defence Related Industries:

> "It is neither affordable nor necessary to strive for complete self-sufficiency in armaments production and all the technologies to support it. However, the South African National Defence Force requires that in certain strategic areas, limited self-sufficiency must be retained and maintained and that in others, the South African National Defence Force needs to remain an informed buyer and user of equipment" [26]

Similar tensions exist in the wider commercial and industrial environment. Organisations must balance their need to learn from the mistakes of others against the potential consequences of disclosing information about their own past failures and successes. Sharing the recommendations that emerge from such incidents may result in the loss of competitive advantage that could otherwise be obtained from these insights. These tensions increase as recommendations are passed across geographical and organisational boundaries. Individual operators may see the benefits of sharing their insights with their fellow workers. Management may be less motivated to share those recommendations with commercial rivals. Ultimately, national political and strategic interests can intervene to prevent the exchange of insights from past failures.

By function...

The previous section has identified some of the problems that arise when investigators draft recommendation that must be applied by colleagues who are not part of the working group that reported an incident. It can be difficult to draft generic remedies that can be applied by groups in other areas. A lack of specific details can remove the directness that characterises many local incident reports. In consequence, particular recommendations can appear to be impositions from external agencies that cannot easily inform the daily working lives of their recipients. These problems that complicate the exchange of information within national boundaries are further exacerbated by the cultural differences that exist between the partners of international systems. These issues can be partly resolved if investigators ensure that recommendations are drafted to target specific functional issues. Less emphasis is placed on generalise from a particular incident so that it can inform a wide range of tasks that are performed throughout an organisation. Greater emphasis is placed on ensuring that similar incidents do not affect the future performance of *the particular task* that was affected by a previous incident.

At the lowest level, task based recommendations can be drafted to support particular working groups within particular units or factories. For example, the following excerpts are taken from the Picatinny Arsenal newsletter published for technicians working on the US 155mm M109A6 self-propelled howitzer, known as the Paladin:

> "An inoperable drivers hatch stop inhibits the ability to properly secure the drivers hatch cover, forces non-operational vehicle status as prescribed in Paladins Operators

> Manual (TM 9-2350-314-10, Feb 1999, Page 2-70, Item 77), and could cause injury if not corrected. Yet, some Paladin personnel improperly use unauthorised field fixes to correct the problem and by doing so promote a potentially dangerous situation. Typically, problems begin when the Grooved, Headless, Pin (NSN 5315-00- 584-1731) breaks (usually when the hatch cover is inadvertently swung open with more than necessary force) then, rather than correcting the problem with authorised parts, a makeshift solution is applied to connect the hatch stop to its shaft. Poorly fitted cotter pins, nails, and similar devices, have been used in place of the Fan Impeller. After continued use, damage to the hatch stop usually occurs causing the hatch stop assembly to become totally inoperable. Units using a Lessons Learned approach to the problem generally maintained a small number of the inexpensive Grooved Pins on hand ($1.78, April 00 Fedlog). When pins were damaged, broken, or became loose, they were quickly replaced. This practice precluded unnecessary damage and replacement of parts, but more importantly a higher degree of operational safety was maintained." [802]

These recommendations illustrate a number of important strengths that can be derived from task-based, local incident reporting. For example, this guidance assumes a high degree of common understanding about the nature of the systems being maintained. Although reference is made to the operators manual and part identifiers, a range of technical terms such as 'fan impeller' or 'grooved pins' can be used without further elaboration. These recommendations dispense with the additional contextual information that is necessary for recommendations that have a wider scope beyond local working groups. Similarly, there is little need to expand on the details of previous violations. It is sufficient to summarise the 'makeshift solutions' for the readers to understand the nature of the incidents that are being addressed. The previous quotation also illustrates some of the weaknesses that limit the utility of such task-based approaches to incident reporting. In particular, the recommendations tend to focus on 'cheap fixes' rather than the large scale investment that may be required to address more systemic failures. Elsewhere, we have reviewed the way in which many aviation and medical incident reporting systems will repeatedly remind staff to 'do better' rather than invest resources in addressing the conditions that led to particular failures [409]. The tendency to rely upon short-term measures is particularly apparent when recommendations are targeted on the tasks or activities performed by individual groups of workers.

We can define a task to be the activities that are required to achieve a particular set of goals [686]. The previous quotations, therefore, examined a very specific and detailed task from the perspective of US Army Engineers at the Picatinny Arsenal. This task focussed approach can also be applied to national and internation objectives. When failures occur at this level, the proposed remedies tend to avoid the short-term solutions that typify more local initiatives. This can be illustrated by the NATO recommendations that were compiled from detailed incident reports and interviews with the personnel who contributed to the peace-keeping missions in Somalia:

> "The evaluation noted many troop contributors' complaint that they were not sufficiently consulted during the formulation stage of the mandate and, thus, had varying perceptions and interpretations during its execution. Many participants in the exercise considered that the original UNOSOM mandate was formulated on political, humanitarian and military assessments, and was prepared, using insufficient information, by officers borrowed for short periods from Member Governments and other peace-keeping operations. Some participants observed that although it was well known that a crisis was unfolding in Somalia, its seriousness and magnitude in humanitarian terms were not fully appreciated. " [624]

It is important to emphasise that even though international organisations may take a more system view of the causes of particular incidents, there is no guarantee that they will be able to solve the problems that complicate high-level tasks such as peace-keeping. Sadly, this point is reinforced by NATO's Department of Peace-keeping Operations review of the Rwanda missions. It was argued that many of the problems and incidents report by NATO forces stemmed from a 'fundamental misunderstanding' of the nature of the conflict [625]. Analysis of individual incidents raised concerns that 'the internal political conflicts within the Government of Rwanda, and the mounting evidence of politically motivated assassinations and human rights violations in the country, were ignored or not explored'.

Previous sections have argued that task focussed recommendations often lead to short-term fixes that ignore more systemic problems. The United Nations recommendations have, however, provided a counter-example. This apparent contradiction can be explained by the very different nature of the tasks that we have considered. Clearly, there are considerable differences between maintaining a driver's hatch and coordinating international peace-keeping operations. It might, therefore, be argued that it is the combination of task-focussed recommendations within a local reporting system that tend to lead to short-term remedies. Task-focussed recommendations at a national or international level are less susceptible to this problem. Previous studies have shown, however, that large-scale systems are far from immune from this problem [409]. There is still an understandable tendency to recommend improved training rather than reassess acquisitions policy.

A number of further limitations affect recommendations that reflect this task-focussed approach to incident reporting. In particular, there is a danger that investigators will fail to consider the importance of particular incidents within the context of a larger operation or production process. For instance, the previous recommendation does not consider the possible acquisition or training problems that led staff to adopt 'makeshift solutions' in the first place. An alternative approach is to embed task-specific recommendations within longitudinal accounts of particular operations. This approach weaves together the findings from a number of different incidents. Any individual may only be directly involved in a small number of the tasks that contribute to the overall operation. However, this longitudinal approach enables them to see how incidents that occur earlier in a process have 'knock on' effects for their own tasks. It also demonstrates that the effects that potential failures in their tasks can have upon the subsequent activities of their colleagues. This 'process-based' approach can be illustrated by the US Army's Engineering Groups analysis of bridging operations. This draws together diverse recommendations from many different stages in a particular bridging operation. Initially, a small 'S3' group compared the tools that the 1st Cavalry Division and the 937th Engineer Group would need to plan and control the operation. This planning exercise identified a number of limitations with current synchronisation techniques and new tools were developed based on 'off-the-shelf' software. Task focussed approaches to incident reporting might have simply presented these recommendations as isolated guidance on the synchronisation of river crossings. This approach is widely adopted in other areas of the US military [802]. In contrast, the engineers of the 937th extended their analysis to integrate it with the recommendations that emerged from the subsequent execution of their plans with the 1st Cavalry Division. Although the tools enabled the engineers to calculate crossing times and schedules for both rafting and bridging, the eventual joint plans did not adequately address some of the fundamental problems that exacerbate the execution of such crossings:

> "A bypassed Orangeland special-forces team on the near shore observed and directed accurate artillery and close air support to destroy the bridge. During the after-action review, it was determined that the critical friendly zones had not been set properly and that the high-to-medium-altitude air defence coverage was inadequate. This action demonstrated that clearing the near and far-shore lodgements is a tenuous and difficult task. One lone member of an opposing force with a radio is the most dangerous person in the crossing area. In an effort to take advantage of the surprise created by the virtually unopposed crossing at Kaw, the division accepted risk by not absolutely ensuring that the crossing site was secure from observed indirect fire before beginning bridging operations. This allowed the division to quickly cross two mechanised task forces but left the ribbon bridge at Kaw vulnerable." [463].

Chapter 3 introduced Perrow's argument that technological failures are unavoidable given that designers are forming increasingly complex interconnections between component systems. For this it follows that even if one organisation has implemented a particular recommendation, there is no guarantee that others will have met the same requirements. This has important implications because failures from one area of a system can propagate through an application to effect later processes that might, themselves, meet the most stringent safety standards. This process-based approach to the presentation of recommendations, therefore, provides eloquent reminders of the mutual dependencies that exist between the component tasks of complex systems. It is important to note, however, that there is still a strong functional bias to the engineer's recommendations. They focus on the particular challenges of the bridging operation and are written from the perspective of those who are tasked to construct and maintain the river crossing. The recommendations do not address the wider strategic significance of the crossings within an exercise as a whole. For example, there is only a cursory description of the problems that the other units in the 1st Cavalry Division faced in exploiting the opportunities, or addressing the threats, that were created at the various crossing points. In other words, the recommendations reflect the functional preoccupations of the engineering group. They propose solutions to incidents that jeopardised their particular tasks. Incidents that occurred elsewhere on the battlefield are not considered. Again, it is important to stress that this example represents a far more general trend. For example, the US Army maintains a number of incident reporting systems that form part of 'Lessons Learned' initiatives. These are organised along functional lines. In addition to the general Center for Army Lessons Learned there is a Center for Engineer Lessons Learned. As we have seen, there is a Contracting Lessons Learned Centre (http://acqnet.sarda.army.mil/acqinfo/lsnlrn/index.htm) and a Medical Lessons Learned unit . The Marine Corps also reflect these functional distinctions by operating separate systems for their combat personnel and for their maintenance staff. The obvious criticism to make of these systems is that there may be important lessons that cross functional boundaries. Some of these are address by Joint Center for Lessons Learned which covers joint forces operations. Other issues that cut across functional boundaries are captured by the US Army Safety Centre. This publishes the safety notices that have been mentioned in previous chapters. It should be stressed that the objectives of these different systems are quite different. For example, the reporting systems maintained by the Safety Centre do not issue the sorts of functional recommendations, for instance on fording tactics, that might appear in the systems operated by the combat engineers. It is important to stress, however, that safety-related incidents and recommendations appear in all of these systems.

This section has argued that the process of drafting effective recommendations is complicated by the geographical scope, the timing and the functional focus of the proposed remedies. Some incidents provide universal insights that can be applied across many different workgroups in particular geographical regions. The Singaporean guidelines on heat injury provide an example of this form of recommendation. Other remedies, such as the proper insertion procedures for the Paladin hatch pin, relate to specific workers performing specific tasks in a few locations throughout the world. It should be stressed that these are not the only distinctions that complicate the drafting of recommendations from incident reports. For example, there are some notable situations in which potential remedies for previous incidents will not be acceptable to both genders. This is illustrated by the guidance provided in the US Army's 'Female Soldier Readiness' [845]. These distinctions have important consequences and investigators must carefully consider their impact on any potential recommendations. For example, a mailshot about the dangers of heat exhaustion may have limited benefits for US Army personnel working at Fort Wainwright in Alaska. Similarly, information about the Paladin hatch mechanisms is of little interest to combat engineers engaged in bridge construction. These geographical and functional distinctions have a profound impact upon many reporting systems. For example, the following list summarises the current titles published in the US Army Safety Centre's leadership guides. As can be seen, some documents provide recommendations that apply to particular geographical regions, including Southwest Asia, Korea, Iraq and the Caribbean. Others relate to particular functions, such as force protection and civilian work force management:

- Annual Training Leader's Safety Guide
- Back Injury Prevention Leader's Safety Guide
- Caribbean Risk Management Leader's Guide
- Civilian Work Force Leader's Safety Guide
- Desert Shield Leader's Safety Guide

- Force Protection Leader's Guide
- Korea Leaders Safety Guide
- Operation Provide Promise Risk Management Leader's Guide
- Operation Support Hope Risk Management Leader's Guide
- Redeployment & Port Operations Leader's Safety Guide
- Southwest Asia Leaders Safety Guide

This list reflects one of the simplest ways in which information can be structured so that users can identify which recommendations are most relevant to their everyday tasks in a particular working environment. Many organisations have developer far more elaborate means of collating and disseminate recommendations. For instance, previous paragraphs have briefly described the US Army's plethora of 'lessons learned systems'. This mixed approach to the dissemination and implementation of recommendations will be the focus of Chapter 14. For not it is sufficient to realise that these diverse information sources provide means of tailoring the presentation and dissemination of recommendations so that they support particular user groups. Unfortunately, they can also create artificial barriers that prevent the free exchange of information about similar incidents between different groups in the same organisation. If any reader believes that these problems are unique to the US Army, it is worth considering the functional and geographical distinctions that are appearing within many healthcare systems. In the UK, different medical specialisms are developing their own guidelines on incident reporting. The Royal College of Anaesthetists' are, arguably, the more well known [715]. At the same time, individual hospitals, NHS trusts and national schemes are all being developed in parallel. This is hardly a situation designed to inspire confidence in the free exchange of information across different organisational boundaries [633].

12.1.3 Conflicting Recommendations

Previous sections have argued that the task of identifying appropriate recommendations is complicated because any remedies may have support a range of tasks that are performed by different operators in many geographical regions. It is further complicated by the ways in which working practices, procedures and technological systems will change over time. Recommendations may, therefore, have to be continually updated if they are to continue to support the safe operation of complex applications. The task of drafting recommendations is further complicated by potential disagreements between investigators, safety managers and national organisations. It is possible to identify at least three different forms for such potential conflict. Firstly, investigators may disagree about the remedies that are appropriate for superficially similar incidents. Secondly, investigators and their managers may disagree over the recommendations that emerge from a particular incident. Finally, safety managers can disagree over the interpretation of particular recommendations. The following paragraphs describe these problems in further detail and provide case studies to illustrate their impact on a number of incident reporting systems.

Different Recommendations from Similar Incidents

Different recommendations are often proposed for incidents that have strong apparent similarities [409]. For example, the US Army's 'Countermeasure' provides military personnel with feedback about a range of safety related incidents. The following incidents appeared in successive numbers of this journal. Both describe two fatalities that resulted from tank drivers using excessive speed during hazardous maneuvers. In the first incident, rapid lateral momentum over a steep slope helped to overturn a seventy ton M1A1. The recommendations paid particular attention to the position of the crew during this incident:

> "Once again, human error became a contributing factor in the loss of a soldier. Leaders must ensure that they and their crew members are positioned correctly in their

> vehicles and are taking advantage of all safety features. The nametag defilade position increases your ability to lower yourself safely inside the vehicle and prevents excessive exposure of body parts to the elements outside. Seatbelts (if provided), guards, clothing and securing equipment enhance your survivability if your vehicle should happen to invert or strike a solid object." [816]

In the second incident, the driver of an M551A1 inadvertently drove into their excavated fighting position so that their vehicle also overturned. Again, the fatality resulted from crust injuries sustained by a soldier who was standing in the hatch above the nametag defilade in the vehicle. In contrast to the previous incident, however, the recommendations did not address the US Army's requirement that all personnel must assume a correct, secured position within any combat vehicles [809].

This apparent difference between the recommendations from two similar incidents can be explained in a number of ways. Firstly, although these incidents resulted in similar outcomes and shared several causes there were also important differences. In the former case, the incident occurred during a daytime exercise. In the second case, the personnel were operating using night vision devices. The recommendations, therefore, focused on the additional requirements for working with limited visibility rather than the requirement to obey seating regulations for combat vehicles. Such differences are not the only reasons why superficially similar incidents might elicit very different remedies. As we have seen in Chapter 11 there are a host of individual and social biases that can affect the analysis of individual failures. These biases may it likely that different investigators may identify different causes for similar incidents. Such problems are further compounded when recommendations are identified in an ad hoc manner without the support of any shared methodology. Later sections in this chapter will describe techniques that have been specifically developed to reduce such apparent differences between the analysis of similar incidents.

A number of further reasons help to explain why investigators derive different insights from similar incidents. New evidence can encourage analysts to revise previous advice. Investigators may also change their recommendations to focus operator attention on particular causes of subsequent incidents. This provides an important communication tool. Over time, the hope is that the readers of Countermeasure and similar journals will learn to recognise the diverse causes of many safety-critical failures. The changing emphasis of particular recommendations can also reflect changes in particular forms of risk assessment. They may signify a decision to focus more on limiting the consequences of an incident rather than reducing incident frequencies. For example, previous speed-related collisions involving combat vehicles had led the Army Safety Centre to reinforce the importance of enforcing recommended speed limits. Subsequent articles focussed more on protective measures that might mitigate the consequences of any collision if a speed-related incident should occur.

It is, therefore, possible to distinguish between inadvertent and deliberate differences between the recommendations that are derived from superficially similar incidents. Inadvertent differences stem from the managerial problems of ensure consistency between the remedies that are proposed for complex events. Investigators often rely on ad hoc methods and do not share the common techniques that might encourage greater agreement. Later sections in this chapter introduce a number of techniques that are deliberately intended to ensure that similar recommendations are derived from similar incidents. However, the large number of incidents that must be investigated by national and international system make it unlikely that such tools will ever provide an adequate solution to this problem. In consequence, Chapter 15 describes a range of search and retrieval software that can be recruited to improve quality control in this domain.

As we have seen, some investigators may deliberately introduce differences between the recommendations that are intended to resolve similar incidents. These differences may stem from individual or group biases that can compromise the value of any subsequent remedial actions. alternatively, deliberate differences may reflect a policy of gradually exposing operators to the underlying complexity of the causes that characterise many incidents. These differences may also reflect the previous success of a system in addressing some of the causes of similar failures. Conversely, they may reflect an apparent failure to address the causes of an incident. Investigators may subsequently focus attention on mitigating the consequences of particular failures. Whatever the justification, it is important that analysts consider presenting the reasons for such apparent differences. For example, the Countermeasure journal dealing with the second incident, described above, deliberately informs its readers that incidents in the special edition will present the diverse range of causal factors that contribute to 'night vision' incidents. This justifies and explains why different recommendations are made after each incident is described. Unfortunately, such contextual explanations are often omitted so that readers have no means of distinguishing deliberate differences with benign explanations from inadvertent differences or differences that are due to the deliberate bias of particular investigators.

Debate Between Investigators and Higher-Level Administration

Previous paragraphs have argued that it is difficult for investigators to ensure that their recommendations are consistent with those of their colleagues. These problems are exacerbated when analysts may deliberately choose to emphasise certain aspects of an incident in their findings. It is also difficult to under-estimate the problems that arise from the sheer scale of many national and international systems. Analysts must ensure consistency between thousands of different reports.

As we have seen differences can arise between recommendations for similar failures. They can also stem from different interpretations of the same incident. One important potential source of dispute stems from the nature of the recommendation 'process' itself. It should be apparent from the use of the term 'recommendation' that these findings are usually recommended by investigatory organisations to a supervising body. For instance, the findings of US Coast Guard reports are typically passed from an individual investigating officer via the Officer in Charge of Marine Inspection to the Commander of the relevant Coast Guard District. Australian Military Boards of Inquiry present their findings to the Minister of Defence and the Federal Government. This process of recommending corrective actions creates the opportunity for disagreement. The Australian Minister of Defence may reject some of the findings made by a Board of Enquiry. Similarly, the Commander of a Coast Guard district may present his reasons for choosing not to implement the findings of an investigating officer. More elaborate mechanisms are also used to approve the recommendations from accident and incident investigations. For instance, the Investigating Officer's Report into the terrorist actions against USS Cole was endorsed by the Commander of US Naval Forces Central Command, by the Chief of Naval Operations and by the Commander in Chief of the US Atlantic Fleet. They 'must approve findings of fact, opinions and recommendations' [836]. Each of these endorsements occurred in a specified order. The Commander of US Naval Forces Central Command provided the initial endorsement, the Commander in Chief of the US Atlantic Fleet was second and the Chief of Naval Operations was last. Subsequent reviewers could not only comment on the report itself but also on the opinions of their colleagues. Most of the comments supported the findings of the investigation. For example, the Chief of Naval Operations stated that "after carefully considering the investigation and the endorsements, I concur with the conclusion of the Commander in Chief, US Atlantic Fleet, to take no punitive action against the Commanding Officer or any of his crew for this tragedy" [836]. There were, however, some disagreements over particular recommendations. There were also disagreements between the endorsing officers! For example, the Commander in Chief of the US Atlantic Fleet observed that:

> "The Investigating Officer and the First Endorser fault the Commanding officer, USS Cole for deviating from the Force Protection Plan he had submitted to his superiors in the chain of command. The Investigating Officer states that had these measures been activated, the attack 'could possibly' have been prevented. I disagree with this opinion, given that those measures would have been inadequate against attackers who were willing to, and actually did, commit suicide to accomplish their attack. I specifically find that the decisions and actions of the Commanding Officer were reasonable under the circumstances." [837]

Other organisations can be commissioned by the ultimate recipients of incident reports to monitor the recommendations that are proposed. For instance, the United States' General Accounting Office was commissioned by members of the senate to review training related deaths. The resulting analysis was not only critical of the recommendations for improving training safety but also uncovered problems with the basis on which those recommendations were made:

> "Our analysis revealed that six deaths categorised by the services as resulting from natural causes occurred under circumstances that could be related to training activities. These were primarily cardiac arrests that occurred during or shortly after the service members had performed required physical training exercises. A typical example of these was a Marine who died from cardiac arrest after completing a required physical fitness regimen. Although he had just completed 5 pull-ups, 80 sit-ups, and a 3-mile run, his death was not considered to be a training death, but rather was classified as a natural cause death." [286]

The Department of Defence responded, in turn, to defend the processes that had been used to investigate particular incidents and the recommendations that had been derived from them. In the final report, the General Accounting Office continued the dialogue by countering these comments with further points about the need to trace whether those recommendations that were proposed had been effectively monitored within individual units.

The complex nature of many incidents often creates situations in which organisations, such as the Department of Defence and the General Accounting Office, hold opposing views about recommendations to avoid future incidents. These conflicts can be difficult to arbitrate. Regulatory or governmental bodies often cannot resolve the differences that exist between the various parties that are involved in the analysis of safety-critical incidents. This point can be illustrated by the Canadian Army's Lessons Learned Centre investigation into their involvement in the NATO Implementation Force and Stabilization Force in Bosnia-Herzegovina (Operation Palladium). They analysed the individual incident reports that had been received during the initial stages of their involvement and made a systematic response to the recommendations that had emerged. The following quotation illustrates how it can be impossible to comply with the competing recommendations that can be made from the different parties who are involved in the analysis of specific incidents:

> "(Reports from Units)...Many units stated that the standard first aid training package (a holdover from Warrior training) lacks realism and that training should be oriented to treating injuries that would be sustained in combat. Many agreed that IV and morphine training were essential components to this training... " During the six months in theatre, no soldier had to give artificial respiration, treat a fracture or do a Heimlich manoeuvre. However, our soldiers did give first aid to 17 bullet-wound cases, 3 shrapnel-wound cases and 7 minefield cases (foot or leg amputated)." As the threat level dropped for latter rotations, unit comments on the need for IV and morphine training waned, there seems to be much debate on the usefulness and dangers of teaching this subject. All unit medical staff strongly recommended that it not be completed because of the inherent dangers that administering IVs or morphine entails...
>
> *(Army Lessons Learned Centre Observation) ...This issue can only be resolved at the highest levels of command in the Canadian Forces and a balance between operational imperatives and medical caution must be found.*" [129]

This quotation provides a detailed example of how it can be necessary to mediate conflicting recommendations. In this instance, the Army Lessons Learned Centre must arbitrate between operational requests for training in the application of morphene and the unit medical staff's concerns about the dangers of such instruction. This example shows how particular recommendations often form part of a more complex dialogue between investigatory bodies and the organisations who are responsible for implementing safety policy. The previous quotation also demonstrates that the political and organisational context of incident reporting systems has a strong influence on the response to particular recommendations. The Canadian Army's Lessons Learned Centre could not reconcile recommendations to expand the scope of trauma training with the medical advice against such an expansion. The fact that they felt uncomfortable with making a policy decision about this matter provides an eloquent insight into the scope of the reporting system and the role of the Centre within the wider organisation. This is not a criticisms of the unit. It would have been far worse if a particular recommendation had been adopted that compromised the reputation of the system or alienated groups who had contributed to the 'lessons learned' process that is promoted by incident reporting.

It is, perhaps, unsurprising that the Lessons Learned Centre should pass such policy decisions to a higher level of authority.

There can, however, also be disagreement at a governmental level. For example, the UK Defence Select Committee examines the expenditure, administration and policy of the Ministry of Defence on behalf of the House of Commons. As part of this duty, it monitors incidents and accidents within the armed forced. The following quotation is taken from the Defence Committee's report into the UK involvement in Kosovo. The first paragraph expresses the Committee's concern about a number of incidents involving Sea Harrier missile configuration. The second paragraph presents the government's response to the Committee's recommendations. The Committee's request for further monitoring is parried by the Government's observation that the problem is not as bad as had been anticipated:

> *"(Committees' recommendation): The resort to cannibalising front-line aircraft in order to keep up the deployed Sea Harriers' availability is clearly a matter to be taken up by the new joint Task Force Harrier's command. We expect to be kept informed of any continuing incidents of damage to the Sea Harrier's fuselage-mounted missiles. (Paras 153 and 176).*
>
> (Government response 57): The Joint Force Harrier is addressing these issues, and the Committee will be kept informed of developments. The problem of AMRAAM carriage in certain Sea Harrier weapons configurations is the subject of continuing in-service trials work, but trials since the potential problem was first identified, together with a longer period of time carrying the missiles, have shown the damage to be much less than feared, and containable within current stock levels and maintenance routines." [792]

This quotation again illustrates the way in which the response to particular recommendations can provide useful insights into the political and organisational context of many incident reporting systems. In this case the government accepts the Committee's request to be informed of subsequent damage to the fuselage-mounted missiles. This acceptance is, however, placed in the context of continuing work on the platform and of the relatively small number of incidents that have been observed. The quotation, therefore, captures the Committee's inquisitorial role and the Government's concern to counter any comments that might be interpreted as politically damaging.

Previous sections argued that investigators must justify any differences between the recommendations that are drawn from similar incidents. Statutory or governmental bodies might also be required to explain why they support particular recommendations and reject others. This was illustrated by the detailed justifications that the US General Accounting Office provided in their rejection of US Army recommendations for training-related deaths. There are, however, situations in which governmental and regulatory bodies are forced to mediate between conflicting recommendations. The Canadian Army's Lessons Learned Centre could not resolve the apparent contradiction between advice for and against specific training in trauma medication. Under such circumstances, particular recommendations must be referred to a higher policy-making body if the position of the regulatory agency is not to be compromised. Even at the highest levels, however, it is important that governmental organisations explicitly justify their response to particular recommendations. For example, the UK Government accepted the Defence Select Committee's request for further information about incidents involving fuselage mounted missiles. It was also careful to explain its response in terms of the most recent evidence about the frequency of such incidents. These explanatory comments can equally be interpreted as political prudence. This underlines a meta-level point; the response to particular recommendations often provides eloquent insights into the political and organisational context of an incident reporting system.

Correctives and Extensions From Safety Managers

The previous section described how differences arise between investigators and the regulatory or governmental organisations that receive their recommendations. Most of the examples, cited above, focus on high-consequence failures rather than the higher frequency, lower severity incidents that are the focus of this book. Similar differences of opinion can, however, be identified over the recommendations that are derived from these failures. These disputes can often be seen in the correspondence that takes place after a report has been published. For instance, most military incidents and accidents are not directly related to either combat or combat training. A large proportion of work-related injuries stem from slips, trips and falls. Others are related to the road traffic incidents that affect the wider community. For this reason, the Canadian National Defence Forces' Safety Digest reported a number of recommendations that were based on several detailed studies of previous incidents [139]. The main proposition in this summary was that car buyers should balance the fuel economy of a vehicle against potential safety concerns. The report argued that the fatality rate for passenger cars increases by 1.1% for every 100-lb decrease in vehicle weight and that in an accident between a Sport Utility Vehicle (SUV) and a car, the occupants of the car are four times as likely to die. Subsequent editions of the Safety Digest carried dissenting opinions from readers who disagreed with the recommendations drawn from previous incidents:

> "(The report) infers that 'bigger' is 'better' for vehicle safety, encouraging readers to buy large automobiles, SUVs, or trucks by feeding their fears. I don't dispute the fact that the larger the vehicle, the higher the chances of occupant survivability in crashes. However, following (this) logic, our family car should be a multi-wheeled armoured fighting vehicle." [141]

The respondent cited studies in which front-wheel drive vehicles with good quality snow tires had outperformed all-season tire-equipped SUVs. They pointed to the problems of risk homeostasis and of decreased perception of risk in larger vehicles. Finally the correspondent argued that the recommendations from the study of previous incidents should have focussed on motivating 'drivers to be more alert, attentive and polite, to practise defensive driving techniques, and to avoid distractions (such as cell phones) and road rage' [141]. This dialogue illustrates the way in which publications, such as the Safety Digest, can elicit useful correctives to the recommendations that can be drawn from previous incidents. Similar responses have addressed more fundamental misconceptions in safety recommendations. For instance, an article about the lessons learned from previous incidents involving electrical systems provoked correspondence that can be interpreted in one of two ways. Either the original recommendations failed to consider the root causes of those failures, as suggested by the respondent, or the respondent had misunderstood the original recommendations:

> "(the report) may leave the erroneous impression that they have discovered new procedures to prevent these types of accidents. The simple fact is that management, supervisors and employees were in violation of numerous existing rules, regulations and safe work practices. Like so many others, this accident was the result of a chain of events which, if carefully examined, often includes all levels - workers, supervisors and management... I have reviewed thousands of accident reports ranging from minor to serious and yes, some fatalities. The vast majority of these reports identify the employee as the cause of the accident. However, study after study has shown that the root cause of accidents is usually somewhere in the management chain. Unless management creates a safety culture based on risk management and unless supervisors instill this workplace ethos in their workers: 'In my shop everyone works safely, knows and follows the rules, and has the right to stop unsafe acts,' and then enforces this view consistently, we will never break the chain and accidents will continue to occur." [136]

As mentioned this correspondence might indicate that the original recommendations did not take a broad enough view of the causes of previous incidents. If subsequent enquiries concurred with this view then additional actions might be taken to ensure that investigators and safety managers looked beyond the immediate causes of electrical incidents. Alternatively, it might be concluded that the respondent had misunderstood the intention behind the original report. In such circumstances, depending on the nature of the reporting system, actions might be taken to redraft the recommendations so that future misunderstandings might be avoided. The previous response not only illustrates how disagreements can emerge over high-level issues to do with the recommendations in an incident report, it also demonstrates the way in which such feedback can challenge more detailed technical advice. The correspondent challenged 'the recommendation that an electrical cane could have been used to effect rescue' during a particular incident [136]. Untrained personnel must not approach any closer than 3.0 meters for voltages between 425V and 12,000V. The national recommendation for trained personnel is no closer than 0.9 of a metre. The respondent concluded that 'the electrical cane shown (in the report) would clearly not be suitable for untrained personnel and only marginal for trained personnel in such a scenario' [136].

Both road safety and the precautions to be taken following electrical incidents are generic in the sense that they affect a broad range of industries. Incident reporting systems also reveal how particular 'failures' can trigger more specialised debates that relate to particular safety issues. For example, the Canadian National Defence Forces' Safety Digest described a series of incidents that stemmed from the need or desire to directly observe particular forms of explosion. One incident occurred during a basic Engineering Officer training exercise. After a number of demonstrations by a tutor, each student prepared and destroyed a piece of ordnance. A student was injured when a fragment shattered a bunker viewport. The subsequent investigation found that the viewports were constructed using four-ply laminated glass. It was designed to withstand a blast equivalent to the detonation of 100 kg of TNT at 130 metres distance with less than 2% glass loss to the inside of the structure [144]. In this case, the glazing performed as designed. Unfortunately, some of the 2% of glass lost to the inside of the bunker lodged in the eye of a student. The recommendations from this analysis focussed on two areas. The first concerns the use of a sacrificial layer of polycarbonate material on the inside of the glazing, not simply to protect against scratches and damage on the external surface. The material would be 'easy to replace when scratched, discoloured or UV degraded, a nd would provide a failsafe final protection for the viewer's eyes' [144]. The second recommendation focussed on the use of periscopes. The offset of the glass elements prevented fragment impact from translating to the viewing side of the optics; "one type of offset viewblock that is in plentiful supply is NSN 6650-12-171-9741 periscope, tank." [144] These recommendations helped to trigger a more general discussion about the technologies that might help to reduce injuries caused by the use of viewports to observe explosions. One correspondent argued that the introduction of sacrificial layers compromised the utility of viewports in other applications. The increasing thickness of the glass 'precluded observation'. In consequence, they recommended the use of video technology:

> "I have seen technology advance to the point where miniature cameras now can be positioned in strategic locations with minimal exposure to blast and fragment impact. Should a lens suffer a direct hit, the replacement cost would be minimal. Lessons learned involving the Coyote vehicle in Kosovo revealed that crews used their digital video cameras to obtain a colour picture rather than relying on the vehicle's integral observation system with its limited monochrome rendering. Closed-circuit TV or a variation thereof permits easy zooming in from a safe distance. It would also be possible to view the demolition site on a number of screens and to record the process for other purposes, including training, slow-motion analysis, replays, and engineering. I have seen video camera lenses smaller than the tip of a pen (using fibre optics) for underwater or high-risk areas (pipeline)." [145]

Such debates can help to increase confidence in particular recommendations. Dissenting opinions and alternative views can be addressed in subsequent publications either by revising previous recommendations or by rebutting the assertions made in critical commentaries on proposed remedies. There is a danger, however, that the results of such dialogues will be lost in many reporting systems. This would happen if the dissenting opinions were not explicitly considered during any subsequent policy decisions. It can be difficult to ensure that such dialogues are both reconstructed and reviewed before any corrective actions are taken. For instance, there is currently no means of reconstructing the thread of commentaries on previous incidents involving the direct observation of explosions. In consequence, safety managers must manually search previous numbers of the Safety Digest to ensure that they have extracted all relevant information. Search tools are available, however, Chapter 15 will describe how these might be extended with more advanced facilities that support the regeneration of threads of debate following from safety-critical incidents. For now it is sufficient to realise that some organisations have devised procedures and mechanisms that are intended to explicitly introduce such debate into the production of incident reports. For instance, sub-regulation 16(3) of the Australian Navigation (Marine Casualty) Regulations, requires that if a report, or part of a report, relates to a persons affairs to a material extent, the inspector must, if it is reasonable to do so, give that person a copy of the report or the relevant part of the report. Sub-regulation 16(4) provides that such a person may provide written comments or information relating to the report. The net effect of these regulations is to ensure that dissenting opinions are frequently published as an appendix to the recommendations in the investigators' 'official' report.

It is important to mention that these dialogues that are often elicited by particular recommendations not only play a positive role in challenging the proposed remedies for particular types of incident. They can also elicit praise that both motivates the continued operation of an incident reporting system and can encourage others to contribute their concerns. There may also be other more specific safety contributions. For instance, one respondent to the Canadian National Defence Forces' Safety Digest publication expressed 'delight' at a report about explosives safety. They then went on to express their disappointment that there had not been any subsequent articles on explosives incidents in the five months since the report had been published; 'Is the world of ammunition and explosives so safe that there is nothing else to write about?' [146]. The correspondent praises the previous article on the causes of explosives incidents. They are also concerned by the relatively low frequency of reports in this area that are summarised in the Canadian National Defence Forces' Safety Digest. This response, therefore, reflects pro-active attitudes to both the underlying safety issues and to the operation of the reporting system. Although such measured reactions are quite rare, they often indicate that an incident reporting system is in good health. If recommendations are challenged then at the very least there is direct evidence that they are being read by the intended audience. If respondents notice that certain types of incidents are under-represented then this can provide evidence of reporting bias. Such responses can also provide valuable feedback about the mechanisms that are used to publicise those recommendations that are derived from previous incidents.

The Dangers of Ambiguity...

Previous sections have argued that the task of drafting appropriate recommendations is complicated by the various correctives that can be issued to address perceived short-comings in the remedies that are proposed in the aftermath of particular incidents. We have described how investigators often issue different recommendations for similar incidents. Such inconsistencies can be intended. For example new remedies may be proposed if previous recommendations have proved to be ineffective. Differences between recommendations can also be unintended. Investigators may not be aware that an incident forms part of a wider pattern of similar failures. Previous sections have also described how regulatory bodies and higher levels of management issue correctives to the recommendations that are proposed by incident investigators. These correctives may directly contradict particular findings. They may also change the emphasis that it placed on particular remedies. Finally, we have argued that well-run reporting systems often elicit debates about the utility of particular recommendations. Operators and managers may also propose ways in which previous remedies might be extended or tailored to meet changing operation requirements. They may also directly contradict the recommendations that have been proposed to address future failures.

The task of drafting effective recommendations is further complicated by the difficulty of ensuring that they can be clearly understood and acted upon by their intended audience. Chapter 14 will describe a range of paper and computer-based techniques that can be used to support the effective communication of particular recommendations. For now, however, it is important to emphasise that there must be stringent quality control procedures to help ensure that the advice that its presented to operation units is unambiguous. This raises an important issue. We have already argued that recommendations must, typically, be expressed at a high level of abstraction if they are to inform the safety of a wide range of different applications. Unfortunately, this also creates opportunity for ambiguity as individual managers have to interpret those recommendations within the context of their own working environment. Peer review and limited field testing can be used to increase confidence that others can correctly interpret the actions that are necessary to implement particular recommendations. If such additional support is not elicited then there is a danger that specific recommendations will be rejected as inapplicable or, conversely, that generic recommendations will be result in a range of potentially inappropriate remedies. At the very least, scrupulous peer review should help to identify 'gross level' inconsistencies. For instance, the US Army Safety Centre reported an incident in which a soldier fell while attempting to negotiate an 'inverted rope descent' [813]. The subsequent investigation a discrepancies between the recommended practices for the construction and use of the obstacle. For example, previous training related incidents had led to the development of standard FM 21-20. This requires that the obstacle should include a platform at the top of the tower for the instructor and the student. A safety net should also be provided. This standard also requires that the obstacle should be constructed to reflect the Corps of Engineers drawing 28-13-95. Unfortunately, this diagram does not include a safety net or a platform. The incident investigators, therefore, concluded that 'confusion exists concerning the proper design and construction of this obstacle'. Following the incident, the army had to suspend the use of their inverted rope descent obstacles until platforms and safety nets had been provided in accordance with FM 21-20. The 28-13-95 diagram was also revised to remove any potential inconsistency.

The previous incident shows how particular failures often expose inconsistent recommendations. Fortunately, many of these problems can be identified before an incident occurs. For example, the US General Accounting Office was requested to monitor the implementation of recommendations following Army Ranger training incidents [288]. They identified a range of problems, not simply in the implementation of those recommendations but also in the way in which those recommendations had been drafted in the first place. For example, one recommendation required that the Army development 'safety cells' at each of the three Ranger training bases. These were to include individuals who had served long enough at that base to have developed considerable experience in each geographic training area so that they understood the potential impact of weather and other local factors on training safety. Safety cells were also to help officers in charge of training to make go/no go decisions. However, the National Defence Authorisation Act that embodied these provisions did not establish specific criteria on the makeup of a safety cell. The General Accounting Office concluded that the approach chosen by the Army 'represents little change from the safety oversight practice that was in place' at the time of the incidents [288]. They also found more specific failures that relate to the implementation of previous recommendations rather than to potential ambiguity in the proposals themselves. For example, the Army Safety Program recommended that safety inspections are conducted on an annual basis. The Fort Benning Installation Safety Office failed to conduct any inspections of training operations safety at the Brigade or its battalions between March 1993 and March 1996.

Chapter 15 addresses the problems and the benefits of monitoring incident reporting systems. It is important to stress, however, that inspections such as that performed by the US General Accounting Office on Ranger Training, can satisfy several objectives. These inspections can be used to expose deliberate failures to implement particular recommendations. They can identify inadvertent neglect; situations in which staff did not know that particular recommendations had been made. These audits also help to recognise genuine difficulties in the interpretation and implementation of remedial actions. Arguably the most significant benefit of such monitoring is that it can be used to institutionalise procedures that help to ensure compliance with key recommendations. For example, the Ranger investigation found that inspections by the Infantry Center, Brigade, and the Fort Benning Safety Office did not monitor compliance with safety controls. In particular, they failed to check that training officers set up minimum air and land evacuation systems before daily training. They also failed to monitor whether instructors adhered to rules prohibiting deviations from planned swamp training routes. The General Accounting Office report concluded that:

> "The inspections are focused instead on checklists of procedural matters, such as whether accidents are reported and whether files of safety regulations and risk assessments are maintained. If the important corrective actions are to become institutionalised, we believe that formal Army inspections will have to be expanded to include testing or observing to determine whether they are working effectively." [288]

The previous paragraphs have argued that monitoring programs can be used to detect potential ambiguity in the recommendations that are issued by incident investigators. They can also assess whether or not those recommendations are being acted upon. This approach does, however, suffer from a number of limitations. Unfortunately, the US General Accounting Office's review of Ranger training only provide a very limited snapshot of one particular area of activity. It ran from September through November 1998 'in accordance with generally accepted government auditing standards' [288]. It involved briefings from Brigade officials. Inspectors observed training exercises and reviewed safety procedures at each battalion's facilities. To determine the level of compliance, they interviewed Brigade officials. They also reviewed Army and Infantry Center inspection regulations, procedures, and records. Personnel were deployed to the Department of the Army headquarters, Army Infantry Center, Ranger Training Brigade headquarters, and the Ranger training battalions at Fort Benning, Dahlonega, Georgia, and Eglin Air Force Base, Florida. The extensive nature of such investigations helped to improve the quality of the eventual report. It also, however, contributed significantly to the costs associated with ensuring compliance. Such techniques cannot easily be applied to support local incident reporting systems where funds may be very tightly controlled. Conversely, they cannot easily be applied to monitor the implementation of recommendations throughout large-scale national systems. For instance, the Modification Work Order (MWO) program was intended to ensure that safety alerts and other maintenance notices were consistently implemented across the US Army [287]. The objective was the enhance fielded weapon systems and other equipment by correcting 'any identified operational and safety problems'.

The implementation of this program was complicated by the number of advisories that it had to track. For example, the US Army approved 95 Modification Work Orders for its Apache helicopter between 1986 and 1997. The implementation of this program was further complicated by the diverse nature of these recommendations. For example, one procedure introduced a driver's thermal viewer, a battlefield combat identification system, a global positioning receiver and a digital compass system into Bradley Fighting Vehicles. The introduction and integration of such relatively sophisticated equipment poses considerable logistical challenges. The MHW program was also intended to monitor less complex modifications. For example, early versions of the Army's High Mobility Multipurpose Wheeled Vehicles utilised a two-point seatbelt restraint system. This did not contain the inertial stopping device that is a standard feature of most civilian vehicles [814]. In consequence, users must remember to remove all of the slack from the retractor and to tighten the seatbelt. This procedure was described and recommended in a safety advisory (TM 9-2320-280-10). Modification Work Order 9-2320-280-35-2 then recommended the installation of a three-point seatbelt system.

A centralised database was developed to record the progress of different maintenance recommendations. Queries could be issued by Army headquarters officials and Army Materiel Command officials to ensure that individual units met the timescales and objectives that were recommended in safety notices. Unfortunately, the centralised database was discontinued following a structural reorganisation in 1990. Control over modification installation funding was transferred from the headquarters level to the individual program sponsors who are responsible for major weapon systems, such as the Abrams tank, or for product centres that support particular pieces of equipment, such as the Squad Automatic Weapon. The result of this decentralisation was that 'Army headquarters and Army Materiel Command officials do not have an adequate overview of the status of equipment modifications across the force, funding requirements, logistical support requirements, and information needed for deployment decisions' [814].

This lack of information also affected field units. It was difficult for maintenance personnel to known which modifications should have been made to particular items of equipment. Similarly, it was difficult to determine which modifications had actually been made. For instance, depot personnel at Anniston Army Depot, Alabama, had to visually inspect 32 National Guard trucks because they had no way of knowing whether two authorised modifications had been made when the vehicles arrived. The difficulties associated with tracking modification recommendations also had knock-on effects. Engineers did not always receive necessary technical information. A General Accounting Office report described how division maintenance personnel did not receive revisions to the supply parts manual for the fuel subsystem on Apache attack helicopters. The aircraft were then grounded and the maintenance team wasted many hours troubleshooting because the old manual did not provide necessary information about a new fuel transfer valve [287]. The lack of an adequate monitoring system created a number of additional logistical problems. For example, it was difficult for engineers to coordinate the implementation of multiple modifications to individual pieces of equipment. In consequence, the same item might be repeatedly removed from service while multiple modification orders were completed. Maintenance teams did not receive adequate notice of modifications. Some items of equipment did not always work together after modifications. This loss of integration further delayed other maintenance procedures and reduced operational capability. For instance, modified parts were removed from Huey utility helicopters. Non-modified parts were then reinstalled because there were no modified parts in stock when the new parts broke. Such practices further exacerbated the problems that were created when responsibility for the database was distributed from headquarters control. The configuration of equipment was not always accurately portrayed in the database used by the maintenance personnel and Army headquarters officials.

A number of recommendations were made as a result of the General Accounting Office report. These included steps to ensure that program sponsors and supply system personnel supported modification orders by providing appropriate spare parts after the initial order had been implemented. The report also recommended that personnel should update technical information whenever a modification order was being performed. Old spare parts were to be 'promptly' phased out and new items were to be added to the units supply system. One of the ironies of incident reporting is that the Accounting Office does not propose monitoring mechanisms to ensure that its recommendations about monitoring practices are effectively implemented!

This section has shown the difficulties of ensuring that the recipients of particular recommendations can unambiguously determine their meaning. It has also illustrated the technical and logistical problems of ensuring that safety recommendations are implemented in a uniform manner across complex organisations. Companies that lack the technological and financial infrastructure of the US Army are likely to experience even greater problems in ensuring that recommendations are successfully implemented. Chapter 15 will describe a number of tools that can be used to address these problems. In contrast, the following sections present techniques that are intended to help investigators identify the recommendations that are intended to combat future failures.

12.2 Recommendation Techniques

A range of techniques have been proposed to help investigator determine the best means of reducing the likelihood, or of mitigating the consequences, of safety-critical failures. Many of these approaches address the problems that were identified in previous sections. For example, some techniques provide methodological support so that the analysis of similar incidents should yield similar findings. They provide a template for any analysis so that disputes can be mediated by reference to the approved technique. Ambiguity can be resolved by encouraging a consistent interpretation of recommendations that are derived from the approved system. The following paragraphs briefly introduce a number of different approaches. These are used to identify potential recommendations from an explosives incident that took place during a nighttime training exercise. The intention was that two maneuver platoons would lead supporting engineer squads across the line of departure. These elements would be followed by a third maneuver platoon. The two lead platoons were to occupy support-by-fire positions. The engineers and the third maneuver platoon were then to occupy 'hide' positions some twenty-five meters from a breaching obstacle. This was to be a triple-strand concertina wire barricade.

The breach exercise was rehearsed a number of times. There was a daytime walkthrough without weapons, munitions or explosives. This was followed by a 'dry fire' exercise in which the plan was rehearsed with weapons but without munitions or explosives. A team leader and two team members would use 1.5 meter sections of M1A2 Bangalore torpedoe to breach the concertina obstacle. The team leader would then pass elements of the initiation system to the team members. They were to tie in the torpedoes to the detonating cords. The initiation system 'consisted of a ring main (detonating cord about 3 to 4 feet formed into a loop) with two M14 firing systems (approximately 4 feet of time fuse with blasting cap affixed to one end) taped to the ring main' [818]. At the opposite end of the M14 firing systems was an M81 fuse igniter that had been attached before the start of the operation. The intention was that the team leader would give each team member one of the M81 fuse igniters. On his command, they were then to pull their M81 and initiate the charge. The breaching team were then to retreat to their original hiding place. The detonation was to act as a further signal for a marking team to use chemical lights to help the following platoons locate the breach.

The actual exercise began when the breaching team approached the concertina objective. The two team members successfully placed their Bangalore torpedoes on either side of a potential breach site. The leader then handed the initiation system to them so that they could tie-in the Bangalore detonating cord lines. The team leader then handed one of the two M81 igniters to the team member on the left-side of the breach. The team leader departed from the original plan when he placed the second M81 on the ground between the two team members. Instead, he handed a bag containing approximately eight meters of detonating cord and an extra M14 initiation system to the team member on the right-hand side of the intended breach. The team leader then radioed the platoon leader to inform them of his intention to fire the charges.

The left-side team member picked up the M81 fuse igniter that had been left on the ground. He also had the original M81 that had been given to him by the team leader. The right-hand team member held the two M81s from the bag. The team members pulled the M81 fuse igniters on the leader's order 'three, two, one, PULL'. A Battalion S3 (operations, planning, and training officer) observed the burning fuses and the added charge in the bag which had been placed to the right of the Bangalore torpedoes. He asked about the additional charge but did not receive any reply. The demolition team and the S3 then moved back approximately twenty-five meters to separate hiding locations. As intended, the detonation acted as a signal for the marking team and a security team to rush towards the intended site of the breach. A second, larger, detonation occurred some three to five seconds after the first. Both of the approaching teams were caught by the resulting blast. The initial detonation had been caused by the additional charge in the bag that had been handed to the team member on the left of the breach. The second explosion was caused by the Bangalore torpedoes.

Chapters 10 and 11 have introduced a number of analysis techniques that can be used to identify the causal factors from this incident. For instance, ECF charts might be used to reconstruct the flow of events leading to the failure. Counterfactual reasoning can then be applied to distinguish causal from contextual factors. Table 12.2 illustrates the results of such an analysis. This tabular form is based on the ECF summaries shown in Tables 10.16 and 10.17. Only causal factors are shown, contributory factors are omitted for the sake of brevity. As might be anticipated, the results of this analysis are similar to the causal findings produced by the US Army technical Centre for Explosives Safety [818]. The original reports do not, however, state whether any particular analytical techniques were used to support the causal analysis of this incident. The justifications associated with the causal factors in Table 12.2 must, therefore, be inferred from the supporting documentation.

The following paragraphs illustrate a range of techniques that can be used to identify particular recommendations once investigators have conducted an initial causal analysis. As will be apparent, there is a considerable imbalance between the number of techniques that might help to identify the causes of an incident and the number of approaches that support the identification of particular recommendations. A cynical explanation for this might be that there is a far greater interest in diagnosing the causes of managerial failure or human error than there is in divising means of addressing such incidents [408]. Alternatively, it can be argued that the identification of recommendations depends so much on the context of an incident and upon the expertise of the investigator that there is little hope of developing appropriate recommendation techniques. However, ad hoc approaches have resulted in inconsistent recommendations for similar incidents. We have also seen ambiguous guidelines that have contributed to subsequent accidents.

The following pages introduce five distinct types of recommendation technique. These distinctions reflect important differences in the role that the particular approaches play within the reporting system as a whole. Some techniques embody the idea that recommendations are imposed upon those who are to 'blame' for an incident. Other techniques reject this approach and provide more general heuristics that are intended to link recommendations more directly to the products of causal analysis techniques, such as ECF analysis. This opens up the scope of potential recommendations; operator failure and human error are not the focus for any subsequent analysis. Other techniques have built upon this link between recommendations and causal analysis by explicitly specifying what actions

Cause	Justification
The breaching team leader failed to turn in excess demo material to the ammunition supply point.	The incident would not have happened if the bag containing the additional M14 initiation system and detonating cord had been handed in.
Excess demolition material was not tied into the ring charge.	The incident would not have happened if all charges had been detonated together.
Addition of the second charge was not planned, practiced or communicated to the other participants.	The incident might not have occurred if the marking and security teams had been aware of the second charge.
There was no appointed command-directed observer/controller at the breaching site.	The incident might not have occurred if a controller had been monitoring the use of the second charge. They might have intervened to prevent the separate detonation of this material.
Breaching team members failed to question or stop the deviated and unpracticed operation.	The incident might not have occurred if team members had questioned the use of the M14 initiation system and detonating cord in the bag.
Battalion S3 (operations, planning, and training officer) recognised but failed to stop the deviated and unpracticed operation.	The incident might not have occurred if they had intervened more directly when their question about the bag went unanswered.
Marking team leader took up hide position closer than the authorised 50 meters to the breaching site.	The consequences of the incident might have been significant reduced if they had been further from the detonation site.
Marking team leader unable to distinguish between the initial (smaller) detonating cord detonation and the larger Bangalore detonation.	The incident might have been avoided if the marking team leader had been able to recognise that the initial detonation was not large enough to have been the Bangalore torpedoes.

Table 12.2: Causal Summary for Bangalore Torpedo Incident.

should be taken whenever particular causal factors are identified. A further class of techniques exploit accident prevention models to identify potential remedies. For instance, barrier analysis approaches look beyond the 'source' of an incident to analyse the defences that fail to mitigate the consequences of particular failures. Unfortunately, a number of practical problems can complicate these broader approaches. Financial and technical constraints can prevent commercial organisations from implementing all of the recommendations that might prevent the causes of an incident and might provide additional protection against the adverse consequences of those failures. A final group of techniques, therefore, exploits concepts from risk assessment to help identify and prioritise the interventions that might safeguard future operations:

1. *recommendations based on blame or accountability.* These recommendation techniques help investigators to remedy the failings of groups or individuals who are 'at blame' for an incident or accident. The intention is to 'put their house in order'. As we shall see, these recommendation techniques are consistent with legal approaches to accident and incident prevention. Prosecution is perceived to have a deterrent effect on future violations. In consequence, recommendations may include an element of retribution or atonement in addition to any particular actions that are intended to have a more direct effect on the prevention of future failures.

2. *recommendation heuristics.* A second class of recommendation techniques take a broader

view both of the causes of incidents and the potential recommendations that can be used to combat future failures. These techniques draft high-level heuristics that are designed to help investigators derive appropriate remedies from the findings of any causal analysis. They provide guidelines such 'ensure that a recommendation is proposed for each causal factor that has been identified during the previous stages of analysis'. Other heuristics describe appropriate implementation strategies. For instance, it might be recommended that 'an individual or organisation is associated with the implementation of any recommendation'. Unfortunately, such ad hoc heuristics provide few guarantees that individual investigators will propose similar remedies for similar failures. There is a danger that inconsistent recommendations will be made within the same reporting system.

3. *navigational techniques (enumerations, lists and matrices).* Instead of focusing on notions such as retribution or blame, a further class of techniques are specifically intended to improve the consistency of particular recommendations. These approaches often enumerate the interventions that investigators should approve in the aftermath of particular failures. For instance a list of recommendations may be identified for each class of causal factors. One consequence of this is that the utility of these techniques is often determined by the quality of the causal analysis that guides their application.

4. *generic accident prevention models.* It can be difficult to enumerate appropriate recommendations for classes of incidents that are still to occur. The dynamism and complexity of many working environments can prevent investigators from identifying effective interventions from pre-defined lists. In consequence, a further class of techniques provides general guidance about ways of improving the barriers and defences that may have been compromised during an incident. Investigators must then interpret this general information within the specific context of their system in order to draft recommendations that will preserve the future safety of an application process. Accident prevention models, including barrier analysis, have been extended to consider mitigating factors. This is important because investigators can use these extended models not simply to consider ways of addressing the causes of complex failures, they can also use them to consider ways of control the consequences of incidents whose causes cannot be either predicted or eliminated [675, 313].

5. *risk assessment techniques.* A number of problems complicate the application of ad hoc approaches and techniques that rely upon accident prevention models to identify incident recommendations. In particular, they provide little guidance on whether particular recommendations ought to have a higher priority that other potential interventions. This is important given the finite resources that many commercial organisations must allocate to meet any necessary safety improvements. It can be argued that such priority assessments are the concern of the regulatory organisations that approve the implementation of investigators' findings. Such a precise division of responsibilities cannot, however, be sustained in more local systems. In consequence, the closing paragraphs of this section consider ways in which risk assessment techniques can be used to identify the priority of particular recommendations. Subsequent chapters consider the regulatory use of these approaches to monitor the overall performance of incident reporting systems.

The following paragraphs assess the strengths and weaknesses of these different approaches in greater detail. Subsequent sections examine the problems of validating the particular remedies that are identified by such recommendation techniques.

12.2.1 The 'Perfectability' Approach

The simplest recommendation technique is to urge operators to do better in the future. In this view, it can be argued that 'if a system demonstrates its underlying reliability by operating without an incident for a prolonged period of time and given that no physical systems have failed then any subsequent failure must be due to operator error'. Such human failures can be corrected by reminding users of their responsibility for an incident. Changes can be made to training procedures and recommended working practices to help ensure that an incident does not recur. Such recommendations are based on the idea that it is possible to avoid future incidents by perfecting previous human 'errors'. This 'perfectability' approach has numerous advantages beyond its apparent simplicity. For instance, reminders are often the cheapest form of remedial action [479]. One consequence of electronic communication facilities is that there be almost no marginal cost associated with sending safety-related emails to members of staff. Of course, such repeated reminders can impair the effectiveness of a reporting system if staff are alienated by repeated reminders about well-known topics [409]. On the other hand, reminders can also be issued more quickly than almost any other safety recommendation. This offers considerable advantages over the length of time that is typically required to implement the re-design of key system components.

Elements of the perfectability approach can be identified in most military regulations. For instance, punitive actions are often prescribed as appropriate remedies when personnel disregard the permits and mandatory obligations that are imposed upon them:

> "the revocation ... of permits to conduct nuclear activities or hold ionizing radiation sources (and) the removal of inventory of ionizing radiation emitting devices from an organisation; and disciplinary or administrative action under the National Defence Act in the case of Canadian Force members and the application of all available administrative measures in the case of Department of National Defence employees. Director General Nuclear Safety may also recommend criminal prosecution." [132]

Similar injunctions have been drafted to ensure that personnel take necessary safety precautions during more mundane activities. For example, the Canadian military stipulates the circumstances in which individuals must wear safety helmets and goggles when operating snowmobiles. Departures from these regulations can be interpreted as instances of individual negligence or of willful violation [133].

Previous researchers have focussed almost exclusively on the use of punishments to 'perfect' operator behaviour in the aftermath of incidents and accidents [700]. It is important to recognise, however, that many organisations operate more complex systems in which rewards may also be offered for notably good performance during near-miss occurrences. For example, the US Army operates a range of individual awards that recognise notably good performance in avoiding incidents and accidents. These include the Chief of Staff Award for Excellence in Safety Plaque, the United States Army Safety Guardian Award, the Army Aviation Broken Wing Award, the Director of Army Safety Special Award of Excellence Plaque, the United States Army Certificate of Achievement in Safety and the United States Army Certificate of Merit for Safety [797]. Such recognition need take little account of the context that may have created the need for individuals to display such acts of bravery and initiative.

Table 12.3 illustrates how the perfectability approach can be applied to the Bangalore case study. As can be seen, each cause describes a failure on the part of an individual. Recommendations are then drafted to ensure that those individuals learn from their apparent mistakes. For instance, the breaching team leader failed to turn in excess materials during the exercise. They should, therefore, be trained in the importance of following such turn-in procedures. Similarly, the marking team leader failed to distinguish between the smaller initial explosion of the detonating cord and the main charge provided by the Bangalore torpedoes. They should, therefore, receive training that might help them discriminate between such different types of detonation.

Table 12.3 deliberately provides an extreme example of the 'perfectability' approach. It illustrates some of the practical problems that arise during the application of this approach. For example, each recommendation is focussed on a particular individual. They, therefore, do not draw out more general lessons. These recommendations neglect the opportunities that an incident might provide for revising the training of all personnel involved in breaching and marking exercises. Further problems stem from the limited effectiveness that such individual recommendations might have in the aftermath of comparatively serious incidents. It is highly unlikely that the individuals involved in this incident would need to be reminded of their individual shortcomings given the consequences of their 'errors'. Such objections can be addressed by drafting recommendations to 'perfect' the performance of groups rather than individual. For instance, the first recommendation in Table 12.3 might be

Cause	Individual Recommendation
The breaching team leader failed to turn in excess demo material to the ammunition supply point.	The breaching team leader should be reminded of the proper procedures for the turn-in of excess munitions.
The breaching team leader did not ensure that any excess demolition material was tied into the ring charge.	The breaching team leader should be reminded of the proper procedures for the disposal of excess munitions. The use of a 'last-shot' to dispose of excess munitions is a dangerous practice and creates the opportunity for such failures.
The breaching team leader's addition of the second charge was not planned, practiced or communicated to the other participants.	The breaching team leader and the exercise safety officer must be reminded of their responsibility to consider the consequences of and communicate necessary information about any unplanned changes to an exercise.
There was no appointed command-directed observer/controller at the breaching site.	The officer in charge of the exercise should be reminded of the need to appoint observers to intervene during potentially hazardous training operations.
Breaching team members failed to question or stop the deviated and unpracticed operation.	Breaching team members must be reminded of their duty to immediately stop any unsafe life threatening act.
Battalion S3 (operations, planning, and training officer) recognised but failed to stop the deviated and unpracticed operation.	Battalion S3 must receive additional training to ensure that they intervene if similar situations arise in future training exercises.
Marking team leader took up hide position closer than the authorised 50 meters to the breaching site.	The marking team leaders must be reminded to follow the required distance regulations specified in IAW FM 5-250.
Marking team leader was unable to distinguish between the initial (smaller) detonating cord detonation and the larger Bangalore detonation.	The marking team leader must be trained to a point where they can distinguish between such different types of detonation.

Table 12.3: 'Perfectability' Recommendations for the Bangalore Torpedo Incident.

applied to all individuals who perform similar tasks to the breaching team leader; 'all personnel must be reminded of the proper procedure for turning-in excess munitions'. This more general approach leads to further problems. For instance, some reporting systems provide participants with apparently random reminders about particular safety procedures. It can be difficult for individuals to follow the justification for these reminders if they are not kept closely informed of the incidents that motivate safety managers to reinforce these particular guidelines. Chapter 14 will describe techniques that can be used to address these potential problems. For now, however, it is sufficient to emphasise that a host of further problems complicate the application of the perfective approach to drafting incident recommendations. In particular, the perfective approach often relies upon demonstrating that individuals have in some way contravened regulations and procedures that they ought to have followed. This creates problems when an incident is not covered by any applicable regulation. It then becomes difficult to argue that individual operators should have intervened to mitigate a potential failure. One ad hoc solution is to continually redraft procedures in a (probably) futile attempt to codify appropriate behaviour in all possible situation. For instance, the US Army Safety Policies and Procedures for Firing Ammunition for Training, Target Practice and Combat contains a requirement that:

> "Accidents caused by firing or evidence that would indicate that the safety provisions of this regulation are inadequate will be reported by letter. The letter must give all pertinent information on the alleged inadequacy of the regulation" [795]

Further problems affect the punitive measures that are associated with the perfective approach. For example, it can be difficult to know exactly what sanctions can be applied to address particular errors and violations. These measures can be influenced by local practices within particular organisations but they are ultimately governed by legislation. Chief Justice Lamer of the Supreme Court of Canada explained in R. v. Généreux in 1992:

> "The purpose of a separate system of military tribunals is to allow the Armed Forces to deal with matters that pertain directly to the discipline, efficiency and morale of the military. The safety and well-being of Canadians depends considerably on the willingness and readiness of a force of men and women to defend against threats to the nation's security. To maintain the Armed Forces in a state of readiness, the military must be in a position to enforce internal discipline effectively and efficiently. Breaches of military discipline must be dealt with speedily and, frequently, punished more severely than would be the case if a civilian engaged in such conduct. As a result, the military has its own Code of Service Discipline to allow it to meet its particular disciplinary needs. In addition, special service tribunals, rather than the ordinary courts, have been given jurisdiction to punish breaches of the Code of Service Discipline. Recourse to the ordinary criminal courts would, as a general rule, be inadequate to serve the particular disciplinary needs of the military. There is thus a need for separate tribunals to enforce special disciplinary standards in the military." [130]

The Chief Justice refers to the importance of punishing breaches of the Code of Service Discipline in order to preserve the 'safety and well-being' of the Armed Forces. This may seem to be a relatively clear-cut decision. There are, however, situations in which there are legal barriers that prevent the application of the 'perfectability' approach even though organisations might want to impose particular sanctions. For instance, a former Sergeant in the Canadian Army found himself as a defendant in a standing court martial when he refused to receive an anthrax vaccination while deployed in Kuwait [142]. His opposition was described as 'unsafe and hazardous'. The case was, however, stopped when the defence cited the Canadian Charter of Rights and Freedoms. This is one of several similar cases in which individuals have used legal arguments to defend themselves against punitive sanctions. Such defences must be provided because there is a danger that superiors may apply 'perfective' sanctions for personal rather than professional reasons. Article 138 of the Uniform Code of Military Justice, section 938 of title 10, United States Code provides one such defence. This article enables a member of the US Armed Forces to seek redress for grievances against a commanding officer and, if redress is denied, to file a formal complaint against that officer. The Judge Advocate General of the Army will then review and take final action on such 'Article 138' complaints.

In more serious cases, sanctions cannot simply be applied by commanding officers. They must be supported by legal argumentation in court martials. Even here, however, there are checks and balances that prevent the arbitrary application of the 'perfective' approach. For example, two of the six appeals currently recorded by the Canadian Judge Advocate General relate to military personnel challenging sanctions that are imposed following safety-related incidents. Such incidents illustrate the more general, pragmatic problems that make it difficult to identify appropriate recommendations within the 'perfective' approach. Training can be ineffective if it is not supported by practical demonstrations and almost constant reminders of the importance of key safety topics. These constant reminders can alienate staff unless properly motivated by concrete, 'real-world' examples. Conversely, more punitive sanctions can be administered by organisations. These legal sanctions are bounded by the civil law and, in the context of our case study, by military law. The development of human rights legislation and of case law that stresses the importance of performance shaping factors as well as individual violations has helped to 'draw the teeth' of the perfective approach in many application domains.

There are a number of theoretical reasons why the perfective approach offers dubious support for investigators and regulators. Recommendations that are intended to perfect operator behaviour often lead to a vicious cycle in which employers become increasingly frustrated by recurring incidents. Reason terms this process the 'blame cycle'. This cycle is based on the notion that operators exercise free will in the performance of their daily tasks [701]. They are assumed to be free to choose between right and wrong, between error-free and error-prone paths of interaction. Any incidents and accidents that do occur are, therefore, partly the result of voluntary actions on the part of the operator. As we have seen, employers and regulators who adopt the 'perfectability' approach are likely to respond to such failures by reminding individuals of their responsibilities and duties. Retraining may be used to reinforce key safety information. Warnings about the consequences of violation are, typically, reiterated after particular incidents. Unfortunately, these recommendations and remedial actions may not address the underlying causes, or performance shaping factors, that created the context in which an 'error' occurred. In consequence, it is likely that there will be future incidents. When these occur, employers and regulators are increasingly likely to resort to additional sanctions and punishments for what they interpret to be willful violations of publicised procedures. Their response to recurrent incidents can be driven by the 'fundamental attribution error' that we have met several times in previous chapters [701]. This arises describes situations in which we ascribe the failure of others to personal characteristics, such as neglect or incompetence, when in similar circumstances we might justify our own mistakes by pointing to contextual factors, such as the level of automated support or time pressures. If punitive sanctions are introduced then they can have the paradoxical effect of making future incidents more likely. They may increase the level of stress in the workplace or may increase a sense of alienation between the employees and their supervisors. In either case, future incidents are likely unless the underlying causes are addressed and so the cycle continues.

To summarise, the 'perfective' approach drafts recommendations that are intended to avoid any recurrence of particular individual errors. This approach is limited because recommendations often address the causes of catalytic failures rather than the causes of more deep-seated managerial and organisational problems. Reason [701], Hollnagel [361], Perrow [675] and Leape [479] have done much to challenge previous applications of this perfective approach. They draw upon a wealth of evidence to suggest that punitive sanctions, individual retraining and constant reminders may have little long-term effect on the future safety of complex, technological systems. There is, however, a need for balance. Any consideration of the context in which an incident occurs must not obscure individual responsibility for certain adverse occurrences. For example, it is possible for risk preferring individuals to alter their behaviour when responding to particular situations in their working environment [368]. It then becomes difficult to distinguish between situations in which those individuals fail to recognise the potential danger inherent in a particular situation, for example because they did not receive adequate training, and situations in which they deliberately choose to accept higher risks in the face of adequate training. There is, therefore, a tension between the need to recognise the impact of contextual or performance shaping factors and the importance of an operator's responsibility for their actions. Many organisations have drafted guidelines that recognise this tension. For instance, the US Air Force's guidance on Safety Investigations and Reports contains the following advice about the drafting of recommendations:

> "5.10.1.5. Write recommendations that have a definitive closing action. Do not recommend sweeping or general recommendations that cannot be closed by the action agency. Vague recommendations addressing the importance of simply doing ones job properly are also inappropriate. However, recommendations to place CAUTIONS and WARNINGS in Technical Order guidance relating the adverse consequences of not doing ones job properly may be appropriate. Recommendations for specific action such as refresher training, implementing in-process inspections, etc. to ensure job duties are being properly performed may also be appropriate since they are specific, and can be closed." [794]

This reflects the tension that exists between the impact of more recent ideas about the organisational roots of many incidents and the 'perfective' notions of free will and individual responsibility. The USAF guidelines reject the 'perfective' notion that individuals should be encouraged to do their job properly. They do, however, accept that it may be necessary to warn operators about the consequences of failing to do their job properly.

12.2.2 Heuristics

Most incident reporting systems provide only a limited guidance about the techniques that investigators might use to derive conclusions for the results of a causal analysis. The NASA procedures and guidelines (NPG 8621.1) that structured the analysis in Chapter 10 recommend seven different causal analysis techniques. In contrast, they offer no suggestions about techniques that might be used to identify potential remedies once causes have been determined [571]. There are good reasons for this reticence. As has been mentioned, a relatively large number of techniques have been proposed to support causal analysis while only a handful have been developed to help structure the identification of recommendations. Those techniques that have been developed are not widely known and tend only to be applied within particular industries, such as chemical process engineering. This contrasts with a technique such as MORT which has been more widely applied and is known throughout many different safety-critical domains;

There are further reasons why some organisations fail to identify appropriate recommendation techniques. Many organisations have failed to propose specific techniques to support the process of identifying recommendations because there is a natural concern that such an approach might unnecessarily constrain the skill and judgement of investigators. A particularly important issue here is that considerable domain knowledge is needed when identifying appropriate remedies. Such expertise cannot easily be synthesised within recommendation techniques. This can be contrasted with causal analysis where it is possible to identify broad categories of failure that contribute to many different incidents. It is possible to challenge these diverse arguments. For instance, the lack of consistency between the recommendations of many investigators in the same industry seems to demonstrate that many do not currently share the same, necessary level of expertise. Similarly, as we shall see, some recommendation techniques have succeeded in identifying generic remedies that can be applied to particular causes in a broad range of industries.

Finally, management may lack the will or the commitment necessary to ensure that investigators follow approved methods when proposing particular recommendations. As mentioned, incident investigators tend to be highly skilled in primary and secondary investigation. Considerable expertise is required in order to direct the causal analysis of safety-critical incidents. In consequence, investigators yield considerable power and influence within investigatory and regulatory organisations. New techniques, that support either causal analysis or the identification or recommendations, can be perceived as a threat to their existing skills and expertise [686]. Many statutory bodies also fail to perform any quality control over the work of their investigators. This leads to a paradox. Investigatory and regulatory organisation do not follow the standardised working practices that they enforce on others.

Many organisations do provide high-level guidance to their investigators. For instance, the Canadian Army's safety program includes a five step guide to accident and incident investigation [131]. These steps are: visit the accident scene; conduct interviews; gather and record evidence; evaluate the evidence and draw conclusion; make recommendations. The following high-level advice is offered to support the final stage of this process:

> "Recommendations:
>
> 31. Once the cause factors have been identified, the investigator(s) recommend(s) preventive measures be taken based on the findings of the investigation. The basic aims when developing preventive measures are as follows: treat the cause and not the effect; ensure that the measures will enhance and not restrict overall operational effectiveness; ensure preventive measures eliminate or control all causes.
>
> 32. Simply recommending that the individual(s) involved by briefed contributes little. It merely indicates fault finding. If human factors (inaction or action - human error) is a

> cause, revising job procedures, training of all employees doing similar tasks and publicity of the accident, to name a few, would be more meaningful and certainly more productive.
>
> 33. If shortcomings in equipment, facilities or other resources are causes, then modifications, substitution or acquisition would be valid recommendations." [131]

This quotation illustrates the importance of eliminating or controlling all causes. Many organisations, therefore, require that investigators explicitly list the remedies that are proposed next to each cause of the incident. This enables colleagues to ensure that each cause is considered in an eventual report. A number of theoretical objections can be raised to this pragmatic objective. For example, the subjective nature of many causal analysis techniques provides few guarantees that this approach will address all of the causes that might possibly be identified in the aftermath of an incident or accident.

The previous quotation stresses the overall objective of operational efficiency. A number of caveats can also be made about this requirement. For example, the guidance does not provide a clear definition of 'operational efficiency'. In practice, therefore, staff may find particular problems in resolving the conflict that often arises between safety concerns and more efficient operational techniques. Paragraph 32 makes the important point that re-briefing soldiers should not be seen as a recommendation. Previous work has noted the tendency of many incident reporting systems to rely upon issuing dozens of similar warning messages [409]. Such 'remedies' provide cheap fixes and may neglect underlying safety issues. The following paragraphs will refer to this as the 'perfective approach' to issuing recommendations. Other organisations have issued more detailed guidance that is intended to help investigators derive particular recommendations from the findings of a causal analysis. For instance, the US Air Force's involvement in aviation incidents has led to the publication of extensive guidance on incident and accident reporting [794]. The following paragraphs use the USAF guidelines to identify a number of high-level recommendation heuristics.

Heuristic 1: Match Recommendations to Each Causal Factor

The USAF guidelines include the generic requirement that 'all mishap investigations should include recommendations to prevent future mishaps'. Like Canadian Army guidance, investigators are urged to match recommendations to each causal finding although exceptions are permitted if they are explicitly justified. Recommendations can also be made against non-causal findings. For example, an investigation may identify alternative ways in which an incident might have occurred. It is, therefore, important to draft recommendations that address both the causal chain that led to an incident as well as any other potential failures that might also have been identified.

Heuristic 2: Assign action agencies for all recommendations

Investigators must clearly identify an agency that will be responsible for ensuring that a recommendation is implemented. Safety management groups should not routinely be tasked to implement particular remedies. In contrast, investigators should identify those groups that manage the resources that are necessary to implement a recommendation. Investigators should also confirm that they have correctly identified a responsible authority providing that this does not compromise their work, for instance by fueling rumours about the potential recommendations.

Heuristic 3: Recommendations Correct Deficiencies

Rather than requiring that an agency should implement a particular solution, investigators should draft recommendations to correct deficiencies. For example, investigators might avoid proposals to 'move the right engine fire push-button to the right side of the cockpit'. In contrast, it would be better to recommend that 'changes should be made to the engine fire push-buttons to help preclude engine shutdown errors' [794]. This second approach goes beyond a simple instruction and helps to provide the rationale behind a particular recommendation. There are further justification for this heuristic. The time-pressures that affect many incident investigations can often prevent investigators from identifying all of the potential ways in which a problem might be addressed. Investigators may also lack the necessary, detailed, domain knowledge that is shared by particular system operators. They might, therefore, be able to device more optimal solutions to that recommended by an investigator in the immediate aftermath of an incident.

Heuristic 4: Recommendations Support Actions NOT Studies

Investigators should be encouraged to draft recommendations that support particular actions. If there is insufficient information upon which to base those actions then studies can be advocated but only as part of the process of implementing the higher-level recommendation. If investigators simply recommend that a study is conducted then there may be no guarantee that any actions will be based on the findings of such an enquiry. Similarly, if a recommendation refers to tests that are incomplete when the report is sent prepared then investigators must identify potential remedies that are contingent upon the outcome of such studies. These different recommendations must be explained and investigators should make explicit reference to the test. They should also explain the reasons why a report was issued before the analysis was completed.

Heuristic 5: Recommendations follow Implementation Paths

It is important that any recommendations take into account the correct procedures and paths for ensuring that corrective actions are implemented effectively. Part of this requirement can be satisfied by ensuring that the recommendation identified an appropriate implementation agency. There may also be other constraints depending on the nature of the recommendation and the organisation in which the incident occurred. For example, investigating officers who recommend changes to military documentation may be required to initiate those changes themselves. This involves the submission of revision requests by submitting the appropriate forms to the relevant office. For example, the USAF guidelines describe the use of the Technical Order System, or AF Form 847, Recommendation for Change of Publication (Flight Publications), according to AFI 11-215, Flight Manual Procedures 'as applicable' [794].

Heuristic 6: Recommendations Acknowledge Minority Opinions

In multi-party investigations, different investigators can have different degrees of influence on the drafting of recommendations. Problems arise when these 'primary' analysts disagree with the remedies proposed by their colleagues. Alternatively, investigators may hold equal influence but are divided into majority and minority opinions. In such circumstances, it is important that the dissenting opinions are voiced. Majority groups or primary investigators must justify their decision not to recommend certain courses of actions.

The USAF guidelines are unusual. They provide detailed heuristics for the identification of particular recommendations. Those heuristics are relatively informal. No explanation is provided for how they were drafted. The reader is not informed of any validation that might confirm the utility of this guidance. They do, however, reflect the pragmatic concerns that are commonly voiced by incident investigators [850]. The US Army's Army Accident Investigation and Reporting Procedures Handbook contains less detailed advice [806]. It does, however, summarise many of the points made in the equivalent USAF publication:

> "Recommendations. Each finding will be followed by recommendations having the best potential for correcting or eliminating the reasons for the error, material failure, or environmental factor that caused or contributed to the incident. Recommendations will not focus on organisational steps addressing an individuals failure in a particular case. To be effective at preventing incidents in the future, recommendations must be stated in broader terms. The board should not allow the recommendation to be overly influenced by existing budgetary, material, or personnel restrictions. In developing the recommendations, the board should view each recommendation in terms of its potential effectiveness. Each recommendation will be directed at the level of command / leadership having proponency for and is best capable of implementing the actions contained in the recommendation." [806]

As can be seen, there are also similarities between these guidelines and those issued by the Canadian Army. Both emphasise the 'effectiveness' of any recommendations. There are also differences. For instance, the US Army explicitly states that investigators need not be 'overly influenced' by existing budgetary constraints. All three of the organisational guidelines in this section emphasise the importance of directing recommendations at a responsible authority. However, the previous quotation not only stresses the need to identify an appropriate agency, it also stresses the need to specify an appropriate level of command within that organisation.

These guidelines are informal. They gather together ad hoc requirements that are intended to improve the quality of recommendations that are produced in the aftermath of safety-critical incidents. They are 'ad hoc' because they have not been integrated into a systematic method or process. Investigators must endeavour to ensure that they obey these guidelines as they develop individual recommendation. It is important to emphasise, however, that these comments should not be interpreted as overt criticisms. Informal guidelines provide important pragmatic advice that is essential given the relative lack of well-developed methods in this area.

Table 12.4 shows how the US Army Technical Centre for Explosive Safety's recommendations from our case study incident can be mapped onto the causal factors that were identified in Table 12.4. As can be seen, this summary explicitly identifies the responsible agency that was charged with implementing the recommendation. The tabular form also illustrates the relationship between recommendations and causal factors. As can be seen, some causal factors are not explicitly addressed. Similarly, some recommendations are not associated with an implementation agency. Table 12.5 also records a recommendation that was made in the incident report but which cannot easily be associated with any of the particular causal factors that were identified from this incident.

This analysis shows how a simple tabular form can be used, together with Army guidelines, as a form of quality control for the recommendations that are made in incident reports. Investigators might be asked to ensure that a recommendation is associated with each of the causal factors. For example, Table 12.5 does not explicitly denote any recommendation that might have helped to avoid situations in which excess demolition material is not tied into a ring charge. It can be argued that this cause is addressed by the previous entry describing how excess material must be turned in. If this analysis were accepted then Table 12.5 should be revised to explicitly associate this recommendation with both causes. Alternatively, it can be argued that this approach would not provide any 'defence in depth'. If excess munitions were not handed in then there is still a danger that the independent firing of charges might cause the same confusion that led to this incident. Under such circumstances, Table 12.5 should be revised by introducing an additional recommendation specifically addressing the detonation of excess material as part of another charge.

As mentioned, the case study incident report does not identify recommendations for each cause nor does it identify responsible authorities for the implementation and monitoring of each recommendation. It is not surprising that our case study does not conform to the US Army guidelines [806]. The recommendations that are cited in Tables 12.4 and 12.5 were derived from material that was used to publicise the remedies that were advocated in the main report. They were not directly taken from the report itself. The example does, however, illustrate the application of these informal guidelines to assess the recommendations that were publicised in the US Army Technical Centre for Explosive Safety's account of the incident. It can also be argued that many of the principles that are proposed in the army guidelines ought to have been carried forward into the accounts that are used to disseminate information about this failure to other engineers throughout that organisation.

12.2.3 Enumerations and Recommendation Matrices

The heuristics that were introduced in the previous section leave considerable scope for individual investigators. They provide guidance about the general form of particular recommendations, for instance by stressing the importance of identifying appropriate implementation paths. They do not directly help investigators to identify appropriate remedies for particular causal factors. In contrast, enumerated approaches list the possible recommendations that might be made in response to particular incidents. For example, the incident involving the Bangalore Torpedoe was analysed according to the US Army's Accident Investigation and Reporting pamphlet PAM-385-40. This

Cause	Recommendation	Agency
The breaching team leader failed to turn in excess demo material to the ammunition supply point.	Training and safety briefings must present and stress proper procedures for disposal/turn-in of excess munitions and/or explosives. Introduction of left over demolition materials into the last shot has been a long-standing accepted procedure. Such action violates the requirement to turn in all excess explosives.	Training and briefing officers
Excess demolition material was not tied into the ring charge.		
Addition of the second charge was not planned, practiced or communicated to the other participants.		
There was no appointed command-directed observer/controller at the breaching site.		
Breaching team members failed to question or stop the deviated and unpracticed operation.	All personnel must have confidence in their authority to immediately stop any unsafe life threatening act and exercise it accordingly.	All personnel
Battalion S3 (operations, planning, and training officer) recognised but failed to stop the deviated and unpracticed operation.	All personnel must have confidence in their authority to immediately stop any unsafe life threatening act and exercise it accordingly.	All personnel
Marking team leader took up hide position closer than the authorised 50 meters to the breaching site.	Inadequate personnel hide distance approximately 25 meters: Required distance (according to IAW FM 5-250) would have been 100 meters for a missile-proof shelter, 200 meters for a defilade position with overhead cover, 50 meters for Command waiver authorised defilade position.	Unspecified
Marking team leader unable to distinguish between the initial (smaller) detonating cord detonation and the larger Bangalore detonation.		

Table 12.4: Guideline Recommendations for Bangalore Torpedo Incident.

Cause	Recommendation	Agency
	The Phase 1 and Phase 2 walk through exercise, prior to live fire operation, is specifically designed to validate the safe execution for all elements of the live fire exercise. A thorough detailed review of all aspects of the operations should have identified the violation of the 50-meter safe hide distance. Had the marking team and security members been properly distanced from the breaching site, their survivability from injury would have been greatly increased.	All personnel, to include command-directed observer/controllers and safety representative, failed to identify violation of the waiver authorised minimum safe hide separation distance during walk through and dry fire iterations.

Table 12.5: Additional Recommendation for Bangalore Torpedo Incident.

enumerates potential root causes. It also provides a list of recommendations that must be considered when drafting the findings from any investigation:

> "**Code: 01, Key Word/Explanation: Improve school training.**
> The improvement recommended should be directed toward the content or amount of school training needed to correct the accident causing error. For example: a. Provide school training for the person who made the error due to not being school trained. b. Improve the content of a school training program to better cover the task in which the error was made. c. Expand the amount of school training given on the task in which the error was made.
> **Code: 02, Key Word/Explanation: Improve unit training.**
> The improvement recommended should be directed toward the content or amount of unit training needed to correct the accident causing error. For example: a. Provide unit training for the person who made the error due to not being unit trained. b. Improve the content of unit training to better cover the task in which the error was made. c. Expand the amount of unit training given on the task in which the error was made.
> **Code: 03, Key Word/Explanation: Revise procedures for operation under normal or abnormal/emergency conditions.**
> The changes recommended should be directed toward changing existing procedures or including new ones. If the change is to an AR, TM, FM, Soldiers Manual, or other Army publication, tell the date when Department of the Army Form 2028 was submitted.
> **Code: 04, Key Word/Explanation: Ensure personnel are ready to perform.**
> The purpose of this recommendation is to encourage supervisors to make sure that their people are capable of performing a job before making an assignment. They should consider training, experience, physical condition, and psychophysiological state (e.g., fatigue, haste, excessive motivation, overconfidence, effects of alcohol/drugs).
> **Code: 05, Key Word/Explanation: Inform personnel of problems and remedies.**
> This recommendation should be used when it is necessary to relay accident related information to people at unit, installation, major Army Command, or Department of the Army levels.
> **Code: 06, Key Word/Explanation: Positive command action.**
> The purpose of this corrective action is to recommend that the supervisor take action to encourage proper performance and discourage improper performance by his people.
> **Code: 07, Key Word/Explanation: Provide personnel resources required for the job.**
> This recommendation is intended to prevent an accident caused by not enough qualified people being assigned to perform the job safely.
> **Code: 08, Key Word/Explanation: Redesign (or provide) equipment or ma-**

> **teriel.**
> This recommendation is made when equipment or materiel caused or contributed to an accident because: a. The required equipment or materiel was not available. b. The equipment or materiel used was not properly designed.
> **Code: 09, Key Word/Explanation: Improve (or provide) facilities or services.**
> This recommendation is made when facilities or services lead to an accident because a. The required facilities or services were not available. b. The facilities or services used were inadequate.
> **Code: 10, Key Word/Explanation: Improve quality control.**
> This recommendation is directed primarily toward the improvement of training, manufacturing, and maintenance operations where poor quality products (personnel or materiel) have led to accidents.
> **Code: 11, Key Word/Explanation: Perform studies to get solution to root cause.**
> This recommendation should be made when corrective actions cannot be determined without special study. Such studies can range from informal efforts at unit level to highly technical research projects performed by Department of the Army level agencies." [796]

This enumeration illustrates some of the problems that arise when attempting to guide the drafting of recommendations in the aftermath of accidents and incidents. As we have seen, the US Air Force guidelines specifically urge investigators not to draft recommendations that involve additional studies. Heuristic 4 in the previous section was that 'recommendations support actions not studies' [794]. In contrast, the US Army guidance includes code 11 that explicitly covers recommendations to perform studies which can identify solutions to 'root causes'. Such inconsistencies are unsurprising given that very few studies have addressed the problems of deriving appropriate recommendations from the outcome of causal analysis techniques.

Table 12.6 illustrates the way in which the PAM 385-40 guidelines can be applied to the Bangalore Torpedoe case study. As can be seen, each causal factor is addressed by one or more of the recommendations proposed by the army guidance material. PAM 385-40 does not specify the way in which an investigator might identify a particular recommendation for any particular causal factors. This is left to the skill and expertise of the analyst. The specific entries in Table 12.6 must, therefore, be validated by peer review. For this reason, it might also be appropriate to introduce an additional column that explains the reason why a recommendation code was associated with each causal factor. For example, a positive command action might address the unplanned addition of the second change because the "supervisor (would) take action to encourage proper performance and discourage improper performance by his people". This example illustrates the pervasive nature of the 'perfective' approach to incident reporting. The US Army guidelines contain several recommendations that reflect this corrective attitude towards operator involvement in accidents and incidents: 01 (improve school training); 02 (improve unit training); 03 (revise procedures...); 04 (ensure personnel are ready to perform); 05 (inform personnel of problems and remedies) and 06 (positive command action). None of the proposed recommendations addresses the organisational and managerial problems that have been stressed by recent research into the causes of accidents and incidents. Similarly, the proposed recommendations only capture a limited subset of the performance shaping factors that have been considered in previous chapters. These are partially covered by recommendation 04 that encourages supervisors to ensure that their teams are properly trained and in an adequate 'psycho-physiological state'.

Such objections can be addressed by extending the list of proposed recommendations. Additional codes can direct investigator towards recommendations that improve communications between different levels in an organisation or between regulators and line management. Unfortunately, the piecemeal introduction of new recommendation codes raises a number of further questions. For example, previous chapters have argued that the nature of incidents will change over time as new equipment and methods of operation are introduced into complex working environments. This argument has been used to stress the problems of identifying the generic causal factors that drive

Cause	PAM 385-40 Recommendation
The breaching team leader failed to turn in excess demo material to the ammunition supply point.	Code 01 - improve school training. Code 06 - positive command action.
Excess demolition material was not tied into the ring charge.	Code 06 - positive command action. Code 10 - improve quality control.
Addition of the second charge was not planned, practiced or communicated to the other participants.	Code 02 - improve unit training. Code 06 - positive command action.
There was no appointed command-directed observer/controller at the breaching site.	Code 07 - provide personnel resources required for the job.
Breaching team members failed to question or stop the deviated and unpracticed operation.	Code 01 - improve school training. Code 02 - improve unit training.
Battalion S3 (operations, planning, and training officer) recognised but failed to stop the deviated and unpracticed operation.	Code 01 - improve school training. Code 06 - positive command action.
Marking team leader took up hide position closer than the authorised 50 meters to the breaching site.	Code 01 - improve school training. Code 06 - positive command action.
Marking team leader unable to distinguish between the initial (smaller) detonating cord detonation and the larger Bangalore detonation.	Code 01 - improve school training. Code 04 - ensure personnel are ready to perform.

Table 12.6: PAM 385-40 Recommendations for the Bangalore Torpedo Incident.

checklist approaches such as MORT, see Chapter 11. Similar problems arise when investigators attempt to enumerate the recommendations that might be used to address these causal factors. It can be difficult to identify appropriate responses to future incidents. If thee could be predicted with any confidence then safety managers would deploy such remedies pre hoc in order to prevent incidents from occurring in the first place!

A number of further problems complicate the use of enumerations. Lists of approved recommendations can guide investigators towards effective remedies. There is equally a danger that they may bias analysts towards ineffective or even dangerous interventions. Chapter 15 will introduce a number of monitoring techniques that can be used to identify such potential problems. It is important to emphasise, however, that the elements in an enumeration must be carefully validated if they are not to advocate ineffective solutions. These problems are exacerbated by the delays that can arise before the publication of revised recommendation lists. In more ad hoc approaches, individual investigators can tailor their interventions to reflect local conditions and personal observations about effective remedies for particular root causes. Such practices can be constrained when analysts must select recommendations from an enumerated list of approved interventions.

PAM 385-40 enumerates the recommendations that US Army investigators must consider when drafting their reports. As mentioned previously, it does not prescribe which particular remedies should be proposed for particular causal factors. This is both a strength and a weakness of this application of a checklist or enumerated approach. This technique relies upon the skill and insight of the investigator to determine whether or not any of the eleven recommendations can be applied. This provides a degree of flexibility that can be important for organisations that are faced with diverse failures in many different geographical and functional areas. This flexibility creates problems. As we have seen, subjective factors and individual biases might affect an investigator's decision to propose one of these recommendations. Any potential inconsistency is reduced by selecting a remedy from the enumeration. There are, however, no guarantees that any two investigators will agree on the same recommendations from that list for any particular incident. Checklist approaches address these potential problems by providing guidance on which recommendations can be best used to address particular causal factors.

It is important to distinguish between recommendation techniques that simply list proposed remedies and those that provide more direct guidance about when to apply particular remedies. The previous section has illustrated the US Army's use of simple enumerations in PAM 385-40. In contrast, Chapter 11 has introduced the use of more directed approaches. For instance, Table 12.7 reproduces the Classification/Action matrices that from part of the PRISMA causal analysis technique. As can be seen, incidents that involve a failure in knowledge transfer within an organisation might result in recommendations to revise training and coaching practices. Failures that stem from operating procedures are addressed by revising procedures and protocols.

	Organisational Factors				
	External Factors (O-EX)	Knowledge Transfer (OK)	Operating procedures (OP)	Manag. priorities (OM)	Culture (OC)
Inter-departmental communication	X				
Training and coaching		X			
Procedures and protocols			X		
Bottom-up communication				X	
Maximise reflexivity					X

Table 12.7: Example PRISMA Classification/Action Matrix (2) [844]

The approach is more 'directed' than the enumeration presented in the previous section because investigators can identify appropriate recommendations by reading down the column that is associated with each causal factor. Conversely, if other participants in the investigatory process propose a particular recommendation then analysts can read along the rows of the Classification/Action matrix to determine whether this would be consistent with previous findings. Table 12.7 only associated a single recommended action with each causal factors. It is important to stress that this need not be the case in all application domains. For instance, problems involving knowledge transfer might be addressed by revised training procedures and by changes in protocols and procedures. Conversely, there may be situations in which 'cultural factors', such as deliberate violations of procedures, cannot simply be addressed by the 'maximise reflexivity' recommendation proposed in Van Vuuren's Classification/Action matrix. In such circumstance, investigators may not be able to directly read off an appropriate recommendation from such a table. Most of the proponents of this approach confirm this analysis by arguing that these matrices are intended as guidelines that can be broken after careful deliberation rather than rules that should be followed in all circumstances. In consequence, these matrices can only be relied upon to increase the consistency of the recommendations made by investigators. They are unlikely to ensure absolute agreement.

Table 12.7 was originally developed by Van Vuuren to help identify recommendations within Healthcare applications [844]. The precise nature of recommendation tables is determined by the context in which they are applied. For example, the causal factors that are represented as columns in the table must reflect the causal factors that are likely to be identified within a particular application domain. Conversely, the recommendations that form each row of the matrix must capture appropriate remedies for those causes. In terms of our case study, the rows of the matrix can be directly derived from the enumeration provided by PAM 385-40 [796]. Fortunately, the same document also provides an enumeration of potential causal factors. For instance, Table B5 lists 'System inadequacies/readiness shortcomings/root causes'. These can be incorporated into the matrix in a similar fashion to the recommendations that were enumerated in the previous section.

Tables 12.8, 12.9, 12.10 and 12.11 are directly derived from the causal codes and recommendation codes that are given in the US Army's guidance on incident and accident investigation [796]. The crosses represent the only additional information that has been introduced into the matrices. These are used to denote those recommendations that might be made given that particular causal factors have been diagnosed. For example, Table 12.8 shows that if an incident had been caused by inadequate supervision by higher command, investigators might consider recommendations that are intended to ensure that personnel were adequately prepared for the tasks that they were presented with (recommendation code 04). Additional recommendations might be drafted to increase the personnel available in an operation (07), to improve facilities (09) or to improve quality control on maintenance and support services (10). Conversely, Table 12.9 can be used to deduce that recommendations to perform more studies (recommendation code 11) might be proposed if there is evidence of inadequate school training (cause code 05) or inadequate unit training (cause code 06).

Not only can the drafting of recommendation matrices help investigators to move from a causal analysis to the findings of an incident report, they can also help to identify potential flaws in the guidance that is provided to investigators. For instance, Table 12.9 lists the causes that PAM385-40 associated with training failures. These include 'habit interference' (cause code 08). This occurs when 'a person makes an accident causing error because task performance was interfered with either in the way he usually performs similar tasks or the way he performs the same tasks under different operating conditions or with different equipment' [796]. As can be seen from the recommendation matrix, it is difficult to identify one of the approved recommendation codes that might be associated with this potential cause. Improved training, possibly following the principles of Crew Resource Management programmes, might address this problem. There is, however, considerable controversy about the effectiveness of such recommendations [410].

The recommendation matrices that we have derived from the PAM 385-40 codes can be applied to the Bangalore Torpedoe incident. For example, previous sections have argued that the Battalion S3 recognised but failed to question or stop the unrehearsed detonation of the excess munitions. This could have been caused by several factors. For instance, it might be argued that this stemmed from environmental factors such as the timescale available to complete the operation or the difficulty of communicating effectively with personnel during a night-time exercise (cause code 21). In such circumstance, investigators might use Table 12.11 to guide their analysis towards particular recommendations. For instance, investigators might advocate that additional measures be taken to ensure that S3's are prepared, in terms of individual training and safety briefings, to ensure that such departures are prevented from occurring (recommendation code 04). Alternatively, investigators might stress the importance of positive command actions on the part of S3's in similar circumstances (recommendation code 06). It is important to recognise, however, that analysts must continue to exercise their skill and judgement in the application of recommendation matrices. For example, Table 12.11 advocates recommendation to improve facilities (recommendation code 09) and to perform more studies (recommendation code 11) in response to the environmental causes (cause code 21), mentioned above. It is difficult to identify ways in which such measures might help to avoid the recurrence of our case study. Investigators might, therefore, argue that they need not draft recommendations to cover all of the potential remedies that are identified in these matrices. The sufficiency of the proposed solutions can be judged by peer review with other investigators. The proposed remedies will, in most cases, also be assessed by regulators and safety managers when they eventually receive the investigators' report.

It might also be argued that 'failure' was caused by a variant of habit interference. The S3 had become habituated to personnel following the approved plan and so failed to identify that the use of excess munitions departed from the approved procedure. Conversely, departures from approved plans might have become so commonplace that the S3 did not interpret the use of the excess munitions as anything 'out of the ordinary'. This analysis raises a number of problems for our application of the recommendation matrices. The causal taxonomy afforded by PAM 385-40 does

LEADER FAILURE

	Cause 01: Inadequate supervision by higher command.	Cause 02: Inadequate supervision by staff officer.	Cause 03: Inadequate supervision by unit command.	Cause 04: Inadequate supervision by direct supervisor, NCO, platoon leader or instructor.
Recommend. 01: Improve school training				
Recommend. 02: Improve unit training				
Recommend. 03: Revise procedures				
Recommend. 04: Ensure personnel ready	X	X	X	X
Recommend. 05: Inform personnel of problems, remedies				
Recommend. 06: Positive command action			X	X
Recommend. 07: Provide more personnel	X			
Recommend. 08: Improve equipment				
Recommend. 09: Improve facilities	X	X		
Recommend. 10: Improve quality control	X			
Recommend. 11: Perform more studies				

Table 12.8: Recommendation Matrix for Leadership Failures

TRAINING FAILURE

	Cause 05: Inadequate school training.	Cause 06: Inadequate unit/on the job training.	Cause 07: Inadequate experience.	Cause 08: Habit interference
Recommend. 01: Improve school training	X			
Recommend. 02: Improve unit training		X		
Recommend. 03: Revise procedures	X	X	X	
Recommend. 04: Ensure personnel ready				
Recommend. 05: Inform personnel of problems/remedies				
Recommend. 06: Positive command action				
Recommend. 07: Provide more personnel				
Recommend. 08: Improve equipment				
Recommend. 09: Improve facilities				
Recommend. 10: Improve quality control				
Recommend. 11: Perform more studies	X	X		

Table 12.9: Recommendation Matrix for Training Failures

STANDARDS FAILURE

	Cause 09: Inadequate written procedures.	Cause 10: Inadequate facilities.	Cause 11: Inadequate equipment.	Cause 12: Insufficient personnel.	Cause 13: Inadequate quality control.	Cause 14: Inadequate maintenance.
Recommend. 01: Improve school training						
Recommend. 02: Improve unit training						
Recommend. 03: Revise procedures	X					
Recommend. 04: Ensure personnel ready						
Recommend. 05: Inform personnel of problems, remedies	X					
Recommend. 06: Positive command action						
Recommend. 07: Provide more personnel				X		
Recommend. 08: Improve equipment			X			
Recommend. 09: Improve facilities		X				
Recommend. 10: Improve quality control					X	X
Recommend. 11: Perform more studies	X	X				

Table 12.10: Recommendation Matrix for Standards Failures

INDIVIDUAL FAILURE

	Cause 15: Fear, Anger	Cause 16: Complacency	Cause 17: Lack of confidence	Cause 18: Haste, Attitude	Cause 19: Fatigue (self-induced)	Cause 20: Alcohol, drugs, illness	Cause 21: Environment
Recom. 01: Improve school training	X	X	X	X	X	X	
Recom. 02: Improve unit training	X	X	X	X	X	X	
Recom. 03: Revise procedures							
Recom. 04: Ensure personnel ready	X	X	X	X	X	X	X
Recom. 05: Inform personnel of problems, remedies							
Recom. 06: Positive command action	X	X	X	X	X	X	X
Recom. 07: Provide more personnel							
Recom. 08: Improve equipment							
Recom. 09: Improve facilities							
Recom. 10: Improve quality control							X
Recom. 11: Perform more studies							X

Table 12.11: Recommendation Matrix for Individual Failures

not distinguish between these very different causes. In consequence, it can be difficult to identify recommendations that might be used to combat these problems. Further problems arise because even if we could unambiguously ascribe the S3's actions to an habituation error there are no specific recommendations associated with this causal factor. In consequence, investigators are free to identify any remedy that is considered appropriate for such an error.

The allocation of recommendations to causal factors in Tables 12.8, 12.9, 12.10 and 12.11 is arbitrary in the sense that it is based on an initial analysis of PAM 385-40. In practice, additional validation would be required before investigators could use such matrices. As we have mentioned, there can be profound consequences if safety managers propose inappropriate or ineffective remedies for the particular causes of adverse incidents. A particular concern is that we have derived these tables from the US Army's published procedures and guidance documents. There are strong differences between these sources and similar publications that guide civilian forms of incident reporting. For instance, the influence of the 'perfective' approach is arguably greater in systems where military discipline and the chain of command are guiding principles. Having raised this caveat, it is important to acknowledge that the recommendation matrices in Tables 12.8, 12.9, 12.10 and 12.11 are still vulnerable to the criticisms raised by Leape [479] and Reason [701]. The focus on individual error and leadership failures obscures the organisational and managerial factors that have been identified in many previous accidents.

A number of problems complicate the use of navigational techniques that are intended to guide investigators towards particular recommendations from lists of approved interventions. For instance, it can be difficult to predetermine a range of appropriate remedies for incidents that have not yet occurred. In consequence, it is unlikely that investigators will be able to identify potential recommendations for all of the incidents that they might encounter. Similarly, the complex nature of many failures can make it difficult to ensure that approved recommendations address all of the detailed causes of particular incident.

Further problems can arise when approved recommendations do not provide sufficient details for investigators to implement them in the aftermath of a particular incident. For instance, the previous analysis of the Bangalore torpedoe case study identified the following causal factor 'Battalion S3 (operations, planning and training officer) recognised but failed to stop the deviated and unpractice operation'. PAM 385-40 codes can be used to classify this cause. For example, Table 12.11 identifies range of individual categories that might be used. These include a lack of confidence (code 17), undue haste (code 18) or problems with fatigue (code 19). As can be seen, however, these are at a more detailed level than the observation that was derived from the US Army's causal analysis. Investigators must, therefore, extend the initial investigation to ease the mapping between the products of the investigation and the classification provided by PAM 385-40. The same problem occurs in reverse when when the matrix approach is extended to identify 'recommended' intervention techniques. Table 12.11 proposes improved school (code 01) or unit training (code 02). Investigators are also encouraged to draft recommendations that ensure a more positive command action (code 06) or that personnel are ready (code 04). At first sight, this might seem to encourage the consistency that has been advocated in previous sections. Such an impression can be misleading. Even if investigators can agree upon a common recommendation code for a causal classification, there is no guarantee that a high-level remedy such as 'improve unit training' will result in similar interventions at an operational level. There are many different ways in which training might be 'improved' the efficacy of such interventions depends entirely upon which techniques are recommended and whether or not they are successfully implemented at the unit level.

12.2.4 Generic Accident Prevention Models

A number of alternate recommendation techniques explicitly acknowledge the problems in classifying causes and then uses such a classification to identify recommended interventions. These techniques exploit a higher level of abstraction than that embodied within the guidance of PAM 385-40. Investigators are then encouraged to introduce additional 'contextual' details into these abstractions. They are expected to exploit their skill and experience to identify the more detailed interventions that are intended to combat future failures. For instance, Haddon identified ten strategies for accident or incident prevention [299]. These strategies are associated either with the source of the energy that is transferred during an incident or with the barriers that protect the system or with the target that is potentially exposed to the energy release. They, therefore, have close links to the from of barrier analysis that was introduced in Chapter 10 that was developed in Haddon's earlier work [298]. These strategies are mentioned now because they have also been proposed as a high-level framework for the identification of recommendations in incident reports [444].

Energy source.
1. Prevent the buildup of energy, this will help to ensure that the conditions for an unwanted release do not slowly accumulate over time.
2. Modify the qualities of the energy, this will help to ensure that appropriate control measures are identified to help prevent any unwarranted releases.
3. Limit the amount of energy, this will minimise the consequences of any uncontrolled release and may make that release easier to control.
4. Prevent the uncontrolled release of energy.
5. Modify the rate and distribution of released energy, this will help to ensure that any unwanted release is stopped at source as soon as possible.

Barriers.
6. Separate the energy source from the target either in time or space, this helps to mitigate the consequences of any energy release.
7. Use physical barriers to separate the energy source and the target.

Target.
8. Increase the resistance of the target to any potential energy flows.
9. Limit any knock-on or consequent damage following any initial energy loss.
10. Stabilise the situation and initiate repairs as soon as possible in case of compound failures.

As mentioned, the components of Haddon's model have been used to provide a high-level framework that is designed to help investigators identify potential recommendations in the aftermath of incidents and accidents [444]. The particular nature of those recommendations will vary from industry to industry and even from incident to incident. The intention is, therefore, not to explicitly provide an enumeration of potential remedies. In contrast, the components of the model are intended to provide an abstract model of those areas in which an investigator might focus any remedial actions.

Table 12.12 shows how Kjellén's [444] application of Haddon's high-level strategies can be used to structure the identification of recommendations. In this case, we have applied Kjellén's approach to identify potential interventions following the Bangalore Torpedoe incident. This example illustrates both the strengths and the weaknesses of the general approach. Haddon's more general model of accident prevention strategies provides a number of high-level prompts that can guide an initial consideration of potential recommendations. The model is based on the notions of barrier analysis, introduced in Chapter 10, and so it avoids some of the myopia associated with 'perfective' approaches. The focus both on causal factors, such as the build-up of energy, and on mitigating factors, such as the resilience of the target, ensure that investigators do not simply focus on the products of a causal analysis when considering the recommendations for an incident report.

Table 12.12 also illustrates some of the potential problems that can complicate Kejellén's application of Haddon's strategies. Although this approach provides important general guidance, it can be difficult to determine what high-level concepts such as 'prevent the build-up of energy' actually mean in the context of a particular incident. Further problems arise when there are clear conflicts between the potential recommendations that might be derived from Haddon's strategy and the operational objectives that govern particular application domains. For instance, Table 12.12 suggests that the rate and distribution of the energy hazard might be altered by possibly increasing the size of breach that the explosives were used against. This would potentially distribute the forces acting on any particular individual who might be caught in a blast during a training exercise. Any intended

Type of Strategy	Case Study Recommendation
	Hazard/Energy Source
1. Prevent build-up	Avoid use of explosives in night-time exercises.
2. Modify quantities	Limit the power of explosives used in night-time exercises.
3. Limit the amount	Limit the quantity of explosives issued to all personnel in a night-time exercise.
4. Prevent release	Limit the number of detonation devices issued in night-time exercises and have procedures for approving release of additional devices only when needed.
5. Modify rate and distribution	Not applicable - possibly increase size of breach area to distribute force?
	Barriers
6. Separate Source and Target	Ensure marking do not proceed until permission to proceed received from the breaching team.
7. Physical barriers	Prevent any detonation without explicit confirmation from a member of the marking team.
	Vulnerable Target
8. Increase resilience	Ensure marking team carry additional protective equipment.
9. Limit damage	Paramedic teams in immediate vicinity.
10. Rehabilitation/initiate repairs	Depends on type of injury.

Table 12.12: Applying Haddon's Ten Strategies to the Bangalore Torpedoe Case Study

reduction in the severity of an incident would, however, have to be offset against the potential tactical problems of alerting the enemy to a failed attempt on their position. It is also important to remember that mission objectives should not be seen narrowly in terms of the short-term outcome from a particular training exercise:

> "Regardless of the training situation, leaders and soldiers must also understand that training exercises are just that training. Under no circumstances should safety be overlooked to achieve a training objective. It is the safety-oriented process that will assist the unit in achieving the mission successfully. Another accident demonstrates the importance of maintaining focus on the objective safely. The unit was engaged in a challenging river crossing operation when the decision was made to float downstream. Even though current readings had not taken place, a safety boat was not on standby, and an exercise participant was not wearing a flotation device, the squad decided to proceed with the mission anyway. Unfortunately, the rivers current was strong enough that it pulled all the team s elements under an anchored barge. Some of the team members survived, but two of them did not. Again, the mission was part of a training exercise. Now we can look back and think of all actions we could have taken to prevent this unfortunate accident; however, now it is too late for the unfortunate participants. Again, leaders must re-emphasise that when encountering an unsafe situation, the mission must now become safety." [807]

Such complex trade-offs between safety and mission objectives should not be surprising. The opening sections of this chapter argued that they are inevitable given that investigators may recommend changes in current operating practices. The key point is, however, that any potential recommendations that appear to fit well with Haddon's accident prevent strategies must also be carefully validated to ensure that they do not result in unintended consequences that might ultimately increase the likelihood of other incidents.

As mentioned, Haddon's strategies provide a high-level framework that Kjellén has used to guide the identification of potential recommendations following incidents and accidents. Some elements of this approach have been developed more than others. For example, barrier analysis is based around strategies 6 and 7 in Table 12.12. Chapter 10 has already referred to its widespread application as part of many causal analysis techniques. Not only can barrier analysis be used to help identify the failure of protection mechanisms, it can also used to identify potential interventions that might avoid future failures. Before providing an example of barrier analysis as a recommendation technique, it is important to emphasise a number of underlying differences between this approach and the others that have been introduced in previous sections. As we have seen, perfective approaches place sanctions on those individuals and groups who are deemed to be responsible for particular failures. The enumeration and matrix approaches that we have analysed typically focus on identify corrective actions for the causes of incidents and accidents. In contrast, barrier analysis typically helps to identify interventions in the accident 'process' that are intended to eliminate or reduce harmful outcomes. It is possible to object that this approach does not address the root causes that provide the 'starting point' for any failure. On the other hand, barrier analysis is supported by the analysis of causal asymmetry that was introduced in Chapter 11. As we have seen, Hausman has argued that it is infeasible to perform 'backwards reasoning' as a reliable means of identifying particular causes from a set of effects [313]. Perrow confirms this when he argues that it is impossible to anticipate the many different causes of technological failure [675]. It, therefore, makes great sense to attempt to control or mitigate those failures that do occur rather than try to eliminate them entirely.

It is possible to identify a vast range of different barriers that might be recommended in the aftermath of an incident. Physical barriers restrict access to hazardous areas, they constrain the flow of energy from a source towards the target. Organisational barriers, such as permit to work schemes, rely upon procedural mechanisms and surveillance activities to achieve similar ends. Barriers may also be active, in other words they are dependent on the actions of operators or systems, or they may be passive. Passive barriers are inherent within a design and are independent of any initiating actions once they are deployed. They must, however, be monitored in case operational demands erode the protection that they afford to the user. For example, the doors of safety cages can be damaged in order to provide greater access to a working area. Kjellén also argues that barriers can be classified as either technical, organisational or social/individual in nature [444]. He provides a detailed list of such barriers that can provide the basis of a checklist approach to the identification of particular recommendations. Unlike some of the previous enumerations, the intention is not to provide a detailed, exhaustive domain specific list. In contrast, these high-level barriers are domain independent and analysts must again apply their skill and experience to interpret them within the context of a particular incident. For example, the following list builds upon what Kjellén calls social and individual barriers. These are intended to prevent future incidents by changing the 'safety culture' in a working environment:

- 1. education, training and experience of personnel;
- 2. feedback on causes and consequences of previous incidents;
- 3. motivational campaigns, safety meetings and awareness raising initiatives;
- 4. feedback rewards for 'safe' performance and punishments for some violations;
- 5. use of automated and peer monitoring systems to assess safety performance.

As with Haddon's original strategies, investigators must translate these high-level barriers into the specific measures that are recommended in the aftermath of an incident. It is entirely possible that this process of interpretation might result in ineffective or even dangerous proposals. For example, there is no guarantee that a motivational campaign will have any effect upon individual behaviour. Similarly, reward and punishment systems can have negative effects if they alienate staff and create workplace conflict [701]. In consequence, it is also important that investigators consider means of validating the implementation of their recommendations. This is addressed in greater detail in Chapter 15. In contrast, the following list extends the previous analysis to summarise a range of organisational barriers to future incidents [444]. As can be seen, these relate to the procedural mechanisms that are intended to promote the safe operation of application processes:

- 6. ensure sufficient numbers of staff;
- 7. monitor implementation and efficacy of all barriers.
- 8. ensure correct levels of expertise and training;
- 9. provide adequate reference documentation to support training;
- 10. provide adequate documentation for emergency procedures;
- 11. rehearse emergency procedures;
- 12. ensure maintenance is effective and timely;
- 13. exploit a 'permit to work' system if maintenance is itself dangerous;
- 14. ensure adequate exchange of information and staff briefings.

It is possible to identify a number of common features between the elements of this barrier analysis and previous recommendation techniques. For instance, '8. ensure correct levels of expertise and training' is similar to recommendation code 01 'improve school training' and 02 'improve unit training' in PAM 385-40. Similarly, '6. ensure sufficient numbers of staff' is similar to recommendation code 07 'provide personnel resources required for job'. Other organisational barriers have not been proposed by more ad hoc approaches to the enumeration of recommendations. For example, the army schemes that were described in previous sections have had relatively little to say about the maintenance activities addressed by items 12 and 13 in the previous list. The following list again extends Kjellén's application of barrier analysis to summarises a number of technical barriers that might prevent the recurrence of previous incidents. As before, these are intended to provide generic recommendations that might be proposed in the aftermath of many different incidents:

- 15. eliminate or reduce hazards in the design of equipment;
- 16. introduce physical barriers to minimise personnel's exposure to hazard s;
- 17. ensure that personnel wear protective equipment whenever necessary;
- 18. ensure that emergency and first aid equipment is provided;
- 19. design workplace to support operators (noise, ventilation etc);
- 20. minimise the use, transportation and handling of hazardous materials.

As with social and organisational barriers, each of these barriers can satisfy a dual role. They can be used to guide the initial design of a safety critical application. For example, an injunction to 'introduce physical barriers to minimise personnels' exposure to hazards' can be used to guide the development of a design. The products of barrier analysis can also be used to identify recommendations in the aftermath of incidents and accidents. For instance, the same injunction might be proposed as a potential remedy in the aftermath of an incident or accident. Table 12.13 builds on this analysis and uses our extended version Kjellén's barriers to identify potential recommendations from the Bangalore Torpedoe case study.

As can be seen, our application of Barrier Analysis identifies a number of potential recommendations that might be used to inform the drafting of an incident report following the Bangalore Torpedoe case study. Potential remedies are again described at a high level of abstraction and must be refined to include the domain details that characterise this particular incident. For example, a requirement to 'ensure effective use of automated and peer monitoring systems' must be translated into particular procedures that can be implemented within the army's command structure. Further validation would then be required to ensure that the particular steps which were taken in the aftermath of an incident actually satisfied this high-level recommendation. These observations are similar to those that were made about the application of more general models of accident prevention,

Cause	Barrier
The breaching team leader failed to turn in excess demo material to the ammunition supply point.	1. Education, training and experience of personnel. 2. Feedback on causes and consequences of previous incidents.
Excess demolition material was not tied into the ring charge.	15. Eliminate or reduce hazards in the design of equipment.
Addition of the second charge was not planned, practiced or communicated to the other participants.	14. Ensure adequate exchange of information and staff briefings.
There was no appointed command-directed observer/controller at the breaching site.	6. Ensure sufficient numbers of staff. 8. Ensure correct levels of expertise and training;
Breaching team members failed to question or stop the deviated and unpracticed operation.	14. Ensure adequate exchange of information and staff briefings 5. Ensure effective use of automated and peer monitoring systems.
Battalion S3 (operations, planning, and training officer) recognised but failed to stop the deviated and unpracticed operation.	5. Ensure effective use of automated and peer monitoring systems. 14. Ensure adequate exchange of information and staff briefings.
Marking team leader took up hide position closer than the authorised 50 meters to the breaching site.	2. Feedback on causes and consequences of previous incidents. 4. Feedback rewards for ‘safe’ performance and punishments for some violations.
Marking team leader unable to distinguish between the initial (smaller) detonating cord detonation and the larger Bangalore detonation.	1. Education, training and experience of personnel. 8. Ensure correct levels of expertise and training.

Table 12.13: Barrier Analysis of Recommendations from Bangalore Torpedo Incident.

illustrated by Table 12.12. This should not be surprising as both approaches share a common root in Haddon's work on incident causation [298, 299].

There are also more worrying similarities. For example, the recommendations identified in Table 12.13 are similar to many of the interventions that were identified using ad hoc heuristics and enumerations, such as those illustrated in Table 12.6. It can be argued that these similarities perhaps reflect particular properties of our case study. The Bangalore Torpedoe incident does not provide a suitable example to demonstrate the differences between these contrasting recommendation techniques. Alternatively, it can be argued that there are very few differences between the application of accident prevention models and more ad hoc techniques. There is, however, a third explanation. Table 12.13 illustrates some of the problems that can arise when recommendation technique are driven directly by causal analysis. For instance, previous sections have argued that investigators must not only focus on the causes of an incident but also on those barriers and controls that help to mitigate its consequences. Unfortunately, Table 12.13 focuses only on remedies for the causes of the incident. It does not consider the performance of triage and evacuation procedures in the aftermath of the incident. This reflects the balance of detail that was provided in the initial US Army report [818]. In consequence, our analysis does not consider certain recommendations: '10. provide adequate documentation for emergency procedures'; '11. rehearse emergency procedures' or '18. ensure that emergency and first aid equipment is provided'. If the initial causal analysis of the incident had also been extended to include mitigating factors, as was done in Chapter 10 then this would have exposed important differences between the recommendations identified in Table 12.6 and those proposed in Table 12.13.

12.2.5 Risk Assessment Techniques

This section began by describing perfective techniques. These approaches focus almost exclusive on exhortations for operators to 'do better' in order to avoid previous failures. Subsequent sections identified a range of techniques that broadened the scope of this analysis. For example, US Army and Air Force heuristics urge investigators to identify recommendations that address each of the causal factors identified during previous stages in an investigation [794, 796]. This more general approach has become embodied within recommendation matrices. Accident prevention models further broaden the scope of any recommendations that are identified in the aftermath of an incident. Not only do they address individual causal factors, they have also been extended to identify recommendations that are intended to strengthen system defences. These including the mitigating factors that can help to control the adverse consequences of particular failures.

The broadening scope of recommendations is appropriate because it reflects a growing recognition that most incidents involve complex interactions between people, systems and the environment in which they interact [675]. It does, however, create a host of practical problems. In particular, it can be difficult for the recipients of an incident report to determine how best to allocate finite resources to support the implementation of all of the diverse recommendations that might be made by investigators. The Canadian report into Operation Assurance illustrates the scale of this problem [128]. This summarised approximately fifty-one recommendations that were made as a result of incidents that occurred during relief efforts in Rwanda. These recommendations included specific measure to improve training at the highest level within the joint forces:

> "Canadian Forces must 'educate leaders and staffs at the most senior levels in both strategic and operational level doctrine processes'. This should be done as a teaching seminar either prior to or in concert with a major command post or computer assisted exercise. The objectives of the exercise should include education and validation with regards to joint doctrine and validation of Joint Forces Headquarters. Thereafter, education/review must be conducted on a routine basis." [128]

It also included more detailed recommendations, for example about the amount of notice that personnel should be given prior to being deployed in remote locations. Similar observations can be made about detailed investigations into single incidents. Apparently simple incident, such as the misuse of commercial heaters in tents, can generate tens of recommendations that range from improved training of personnel through to changes in the monitoring of standards throughout the chain of command [815].

The diverse nature of many recommendations and the sheer volume of remedial actions that can be proposed in the aftermath of an incident can create considerable problems from the recipients of these reports. Risk assessment techniques provide investigators, regulators and end-users with means of prioritising the recommendations that are are made in the aftermath of an adverse occurrence. Previous sections have noted that the particular details of who performs this prioritisation vary between different industries. For instance, in the Air Traffic Management domain there are well specified procedures that govern the reporting of recommendations by investigators back to safety managers who then prioritise their findings [423]. In local incident reporting systems, for example with UK hospitals, the same individuals may identify and prioritise potential recommendations [119]. This informal process is intended to be highly cost-effective in terms of the resources required to perform the analysis. Although it can be carefully tuned to the local working conditions of the units that operate the incident reporting system, this approach is also open to the many subjective biases that can distort risk assessments [2].

A number of organisations have recognised the key relationship between incident reporting and risk analysis. For example, the US Army Safety Program recently devoted an issue of their Countermeasure magazine to 'Accident Investigation: The Other Side of Risk Management' [807]. Of course, this relationship is more complex than the prioritisation of recommendations. For example, the US Army identifies five stages in a risk management process [805, 798]. These can be summarised as follows:

1. Identify hazards.
   Incident reporting provides important guidance in this initial stage of any risk assessment because it provides information about previous failures. Data can be collated from other operational units both within the same organisation and from national and internation groups operating similar processes. There is, however, no guarantee that previous incidents will provide good information about future failures involving novel production techniques.

2. Assess hazards.
   This second stage of risk management is intended to assess the impact of each hazard in terms of potential loss and cost based on probability and severity. More will be said about this in the following paragraphs. However, for now it is sufficient to observe that incident reporting systems not only provide information about previous types of hazard, they can also be used to identify the likely consequences of future failures based upon previous outcomes.

3. Develop controls and make risk decision.
   As control measures are developed, risks must be re-evaluated until they are reduced to a level that is 'as low as reasonably practicable'. This ALARP principle is controversial because it implies that it is possible to identify situations in which the perceived benefits of reducing the risk of a particular failure any further are outweighed by the potential costs of implementing such a risk reduction. Chapter 15 will describe how incident reporting systems can be used to support this aspect of risk management. In theory, it should be possible to demonstrate the effectiveness of particular recommendations by monitoring falls in the frequency and severity of future incidents. This is not always possible given the problems of ensuring the uniform implementation of recommendations and the relatively low frequency of many safety-critical incidents.

4. Implement controls.
   The fourth stage in any risk management program is to implement the controls that are intended to achieve the intended risk reduction. Again incident reporting systems provide an important source of information on the potential benefits of particular forms of control or barrier. Data from other plants can be sued to determine whether the introduction of these measures can create the opportunity for further types of incident, for instance during the installation and 'burn-in' of new equipment.

5. Supervise and evaluate.
   Finally, it may be necessary to monitor not simply the performance of any recommended controls but also to ensure that personnel continue to follow recommended practices. Similarly, it is important to ensure that recommended safety equipment is maintained and operated in a manner that is intended to ensure that it is availability on demand. It is important to remember the problems associated with gathering reliable evidence for violations from incident reporting systems. It can be difficult to obtain evidence about conformance to recommendations, especially if they have been implemented following previous incidents.

As can be seen, several dependencies exist between risk management and the identification of recommendations in the aftermath of an incident. Risk management techniques can be used to determine the relative priority of particular recommendations. If a particular hazard is thought to be very unlikely or if it implies only marginal consequences for the safe operation of an application then the recommendation may be assigned a relatively low level of priority. Conversely, if an associated hazard is predicted to have a high frequency or a relatively large impact on safe and successful production then recommendations to address that hazard will be assigned a high-level of priority. Incident reporting systems can then be used to assess the efficacy of those recommendations that are rated particularly highly using such risk management techniques. If similar incidents continue to occur then the effectiveness of a recommendation may be questioned. Conversely, if incidents occur from hazards that were assigned a low relative priority then the effectiveness of the risk management system can be questioned.

In order to understand the role of risk management techniques in the identification and prioritisation of recommendations it is first necessary to describe the underlying components of a risk management system in greater detail. The fundamental concept behind this approach to the development of safety-critical systems is that:

$$\mathit{Risk} \ = \ \mathit{Frequency} \ \times \ \mathit{Cost}$$

This formula provides a means of assessing the potential effectiveness of any recommendation in terms of reductions in the costs or consequences of an incident. It can also be used to prioritise recommendations in the aftermath of an incident. As mentioned in previous sections, US Army [796] and Air Force [794] guidelines argue that each recommendation must be clearly associated with the results of a causal analysis. The US Army defines a hazard to be ‘any real or potential condition that can cause injury, illness, or death of personnel or damage to or loss of equipment, property or mission degradation’ [805]. In consequence, the same frequency and consequence that are associated with a hazard can also be associated with the causes of an incident or accident. This can inform the identification of recommendations in one of two ways:

1. if recommendations have already been identified using one of the techniques described in previous sections then risk assessment techniques can be used to identify the priority of each proposed remedy in terms of the risk associated with the cause that it is intended to address;

2. if recommendations have not already been identified then a risk assessment can be performed for each cause. The results of this analysis help to establish a partial ordering that can be usd to allocate finite investigatory resources. Greatest attention should be paid to finding appropriate recommendations for those causes that are assumed to pose the greatest continuing threat to a system.

Unfortunately, a number of factors complicate the application of the previous formula to guide the prioritisation of recommendations. The previous formula is a simplification. Subsequent paragraphs will introduce concepts such as risk exposure that must also be considered when attempting to assess the priority of particular recommendations. If such factors are not taken into account then it is possible to assign relatively low priorities to recommendations that could have a relatively large impact upon the risk of future incidents because investigators fail to accurately assess the potential frequency of a particular causal factor or hazard. It can be argued from the previous formula that the risk associated with a hazard will fall if either its frequency is reduced or the costs associated with that hazard fall. This, of course, assumes that such reductions are not offset by a corresponding fall in the other component of the equation. For instance, Bainbridge has argued that the implementation of many safety recommendations reduces the frequency of a particular cause but can also increase the consequences of those hazards when they do occur [65]. This one of several 'ironies of automation'; the relatively low frequency of certain failures can leave operators unprepared to intervene in adverse incidents.

Further problems stem from attempts to derive numerical estimates for the consequent cost of a particular hazard. The amount of money that must be spent in the aftermath of previous incidents can prove to be an extremely poor indication of what might have to be paid in the future. The increasing use of litigation within certain Healthcare systems has resulted in massive changes in the scale of compensation that must now be paid following many adverse incidents [453, 633]. There are well established mechanisms for calculating the potential liability associated with fatalities and personal injuries. It can, however, be difficult to predict the potential scale of such injuries that might result from future incidents. The costs associated with air collisions can vary greatly depending on the numbers of ground fatalities that are factored into any calculation. It is also difficult to predict the punitive elements that can be introduced during the settlement of claims within some legal systems In consequence, most organisations avoid precise numerical assessments for the potential costs associated with particular hazards. In contrast, they rely upon subjective bands that are described using keywords. This is an approach that complements the use of such terms within HAZOPS [27]. For instance, the US Army encourages risk managers to consider a number of basic categories that can be used to describe the consequences associated with a particular hazard [805]. The costs of an incident are assessed in terms of the expected degree of injury, property damage or other 'mission-impairing' factors:

1. Catastrophic: death or permanent total disability, system loss, major damage, significant property damage, mission failure.

2. Critical: permanent partial disability, temporary total disability in excess of 3 months, major system damage, significant property damage, significant mission degradation.

3. Marginal: minor injury, lost workday accident, minor system damage, minor property damage, some mission degradation.

4. Negligible: first aid or minor medical treatment, minor system impairment, little/no impact on mission accomplishment.

A number of caveats can be raised about the interpretation of these different categories. The relatively low costs associated with near-miss incidents can persuade organisations to underestimate the consequences of a potential accident. It is for this reason that some organisations have argued that the cost component of the risk management equation, given above, should only be calculated in terms of the *worst plausible outcome* of an incident. If pilots were able to narrowly avert a collision then safety managers might assess the costs of such an occurrence in terms of the potential loss of both aircraft. This appears to be a rationale and well considered approach. Problems arise, however, when investigators must determine what 'worst plausible outcome' actually means for specific incidents. This issue was addressed in more detail in Chapter 4.

Chapter 11 notes the difficulty of obtaining quantitative data about incident frequencies. Data from bench trials and experimental observations often cannot be replicated in complex, working environments. Conversely, observational data and information from automated logging systems can be difficult to calibrate and interpret. Incident reporting systems provide a partial solution to these problems. They provide information about 'actual' incidents in 'real' working environments. Forms can be designed to elicit the information that is necessary to interpret observations about adverse events. In confidential systems it is possible to gather additional information about the context in which failures occur. As we have seen, however, participation bias and relatively low submission rates create significant problems for the use of incident reporting data as a 'raw' source for risk management. The US Army [805], therefore, also provides guide-words that describe the frequency of a potential hazard:

1. Occurs often, continuously experienced.
2. Occurs several times.
3. Occurs sporadically.
4. Unlikely, but could occur at some time.
5. Can assume it will not occur.

A number of further problems complicate the application of this approach to risk management. Previous paragraphs briefly mentioned that the risk associated with a particular hazard is partially determined by the length of exposure to that hazard. This must take into account both the cumulative duration of any exposure but also the summative effect of individual exposures across different operational units. These issues were not considered in previous formulae. One means of addressing this omission is to refine the subjective categories that are used to describe the potential frequency of a hazard. This is the approach that is advocated by the US Army's guidance on Risk Management, illustrated in Table 12.14 [798]. One consequence of adopting this approach is that it can introduce additional complexity into the superficial simplicity which is an important strength of the initial frequency definitions.

As mentioned, the previous risk assessment formula can be applied together with the previous definitions of consequence and frequency to estimate the risk that is associated with particular hazards. In practical terms this is accomplished using a risk assessment matrix similar to that presented in Table 12.15. The use of such matrices has important consequences for the use of risk analysis to drive the prioritisation of incident recommendations. As can be seen, Table 12.15 supports the high level classifications of hazards into Extremely high, High, Moderate and Low risks. Such distinctions are unlikely to provide a total ordering over the many different recommendations that are made in the aftermath of safety-related incidents. In consequence, even if analysts do resort to the use of risk assessment techniques to supplement an incident investigation they will still have to exploit a range of additional techniques to rank individual recommendations within these gross categories.

Table 12.15 can be used in conjunction with the US Army guidance on frequency and consequence assessment to prioritise the recommendations that were identified by the investigation into the Bangalore Torpedoe case study. This process begins by performing a risk assessment of the causal factors that were identified in the aftermath of this incident. This approach is justified by the Army guidance that points to the close relationship between the hazards that are considered in any risk assessment and the causes of previous accidents [805]. Table 12.16 illustrates the results of such an analysis. As can be seen, a frequency and criticality level are associated with each of the causal factors. The subjective nature of these assessments makes it important that investigators also document the justification for the allocation of particular levels to each of the causal factors. For instance, the breaching team leader's failure to turn in excess demo material was classified as a likely occurrence on the basis of comments made by the investigating officer: "Introduction of left over demolition materials into the last shot has been a longstanding accepted procedure. Such action violates the requirement to turn in all excess explosives..." [818]. Similarly, the breaching team members' failure to question or stop the deviated and unpracticed operation was assessed as being unlikely. This was based on an analysis of previous exercises in which phase one and phase two walkthroughs established the pattern for an operation and helped personnel to question deviations from the planned actions. Such justifications might be explicitly included within risk assessment documents such as Table 12.16.

Previous paragraphs have briefly described the problems of assessing the likely consequence of a particular hazard in the aftermath of an adverse occurrence. It might be argued that there were no consequences from any of the particular causes of a 'near-miss' incident. In contrast, if we apply the 'worst plausible outcome' assumption then almost every cause can have potentially catastrophic outcomes. This dilemma can be illustrated by assigning a criticality level to the observation that there was 'no appointed command-directed observer/controller at the breaching site'. It is difficult to argue that the lack of an observer led to mission failure, 'death or permanent total disability'

FREQUENT (A) Occurs very often, continuously experienced

Single item	Occurs very often in service life. Expected to occur several times over duration of a specific mission or operation. Always occurs.
Fleet or inventory of items	Occurs continuously during a specific mission or operation, or over a service life.
Individual soldier	Occurs very often in career. Expected to occur several times during mission or operation. Always occurs.
All soldiers exposed	Occurs continuously during a specific mission or operation.

LIKELY (B) Occurs several times.

Single item	Occurs several times in service life. Expected to occur during a specific mission or operation.
Fleet or inventory of items	Occurs at a high rate, but experienced intermittently (regular intervals, generally often,).
Individual soldier	Occurs several times in career. Expected to occur during a specific mission or operation.
All soldiers exposed	Occurs at a high rate, but experienced intermittently.

OCCASIONAL (C) Occurs sporadically.

Single item	Occurs some time in service life. May occur about as often as not during a specific mission or operation.
Fleet or inventory of items	Occurs several times in service life.
Individual soldier	Occurs some time in career. May occur during a specific mission or operation, but not often.
All soldiers exposed	Occurs sporadically (irregularly, sparsely, or sometimes).

SELDOM (D) Remotely possible; could occur at some time.

Single item	Occurs in service life, but only remotely possible. Not expected to occur during a specific mission or operation.
Fleet or inventory of items	Occurs as isolated incidents. Possible to occur some time in service life, but rarely. Usually does not occur.
Individual soldier	Occurs as isolated incident during a career. Remotely possible, but not expected to occur during a specific mission or operation.
All soldiers exposed	Occurs rarely within exposed population as isolated incidents.

UNLIKELY (E) Can assume will not occur, but not impossible.

Single item	Occurrence not impossible, but can assume will almost never occur in service life. Can assume will not occur during a specific mission or operation.
Fleet or inventory of items	Occurs very rarely (almost never or improbable). Incidents may occur over service life.
Individual soldier	Occurrence not impossible, but may assume will not occur in career or during a specific mission or operation.
All soldiers exposed	Occurs very rarely, but not impossible.

Table 12.14: US Army Guidance on Hazard Probability [798].

	A. Frequent	B. Likely	C. Occasional	D. Seldom	E. Unlikely
1. Catastrophic	Extremely high	Extremely high	High	High	Moderate
2. Critical	Extremely high	High	High	Moderate	Low
3. Marginal	High	Moderate	Moderate	Low	Low
4. Negligible	Moderate	Low	Low	Low	Low

Table 12.15: Risk Assessment Matrix.

Cause	Frequency	Consequence	Risk Assessment
The breaching team leader failed to turn in excess demo material to the ammunition supply point.	B. Likely	1. Catastrophic	Extremely high
Excess demolition material was not tied into the ring charge.	D. Seldom	4. Negligible	Low
Addition of the second charge was not planned, practiced or communicated to the other participants.	D. Seldom	1. Catastrophic	High
There was no appointed command-directed observer/controller at the breaching site.	B. Likely	1. Catastrophic	Extremely high
Breaching team members failed to question or stop the deviated and unpracticed operation.	E. Unlikely	4. Negligible	Low
Battalion S3 (operations, planning, and training officer) recognised but failed to stop the deviated and unpracticed operation.	E. Unlikely	1. Catastrophic	Moderate
Marking team leader took up hide position closer than the authorised 50 meters to the breaching site.	B. Likely	2. Critical	High
Marking team leader unable to distinguish between the initial (smaller) detonating cord detonation and the larger Bangalore detonation.	D. Seldom	4. Negligible	Low

Table 12.16: Risk Analysis for Bangalore Torpedo Incident.

unless we know the context in which this hazard occurred. If excess material was being used in an unscheduled procedure then the lack of an observer can have catastrophic consequences. In other contexts the consequences are much less severe. This illustrates the need to provide additional guidance for investigators who must determine the potential future consequences of such causal factors in a variety of different contexts. Ideally, we would like a rule or form of argument that plays a similar role to counterfactual reasoning in many causal analysis techniques [470]. Without such a decision procedure, investigators must continue to rely upon their expertise and judgement when determining the consequence of future hazards. As before, it is important that others can follow the justifications that support such judgements. For example, Table 12.16 assigns negligible consequences to the 'breaching team members failed to question or stop the deviated and unpracticed operation' because this last line of defence should not be relied upon given the stress levels and distractions associated with nighttime operations. Of course, other investigators might argue that the consequences of breaching such a final barrier are critical or catastrophic. The key point here is that by documenting the justifications for such an allocation, it is then possible for other analysts to validate the reasons for prioritising the recommendations that are intended to address particular causes of an incident.

We have argued that the priority of a recommendation can be determined by assessing the risk of the causes that it is intended to address. This depends upon the recognition that the causes of incidents provide valuable information about the hazards that threaten the future operation of safety-critical systems [805]. It is important to emphasise, however, that although all causes can be though of as hazards, it is not the case that all hazards are causes. In particular, there may be potential failures that have not yet contributed to particular incidents. Investigators must consider this issue when assessing the priority of a recommendation. For example, our case study did not involve a friendly fire incident. The introduction of a controller/observer at key positions during a night exercise might also help to reduce the risks associated with this other form of hazard. Hence it can be argued that the priority of this recommendation ought to be increased to reflect the additional perceived benefit to be derived from such an intervention. It is also important to reiterate the argument that was made in the closing sections of Chapter 11. Causal asymmetries imply that there may be a number of different alternative causes for any particular incident. In consequence, investigators must consider the relative important of recommendations that will not simply address the causes of a particular incidents. They must also prioritise recommendations that address alternative causes that might have resulted in the same or similar failures. For instance, there are numerous ways in which the marking party might have suffered similar injuries given that they were too close to the site of the breach when the Bangalore Torpedoes were deployed. A comprehensive risk analysis would, therefore, consider these different causal paths when determining the relative priority of recommendations that might ensure conformance to the distance requirements in the FM 5-250 [818].

The US Army promotes a five stage process of risk analysis: identify hazards; assess hazards; develop controls and make risk decisions; implement controls; supervise and evaluate [798]. Previous paragraphs have described how, in the context of incident reporting, causal analysis techniques can be used to identify the particular hazards that lead to an incident or accident. Hazard assessment techniques can then be used to derive a partial ordering that prioritises those causes. The third step in the process is to identify 'controls and make risk decisions'. This stage can be implemented using the recommendation techniques that have been introduced in this chapter. For example, the US Army's FM 100-14 advocates an approach that has much in common with barrier analysis:

> "After assessing each hazard, leaders develop one or more controls that either eliminate the hazard or reduce the risk (probability and/or severity) of a hazardous incident. When developing controls, they consider the reason for the hazard not just the hazard itself. Controls can take many forms, but fall into three basic categories educational controls, physical controls, and avoidance. Educational controls are based on the knowledge and skills of the units and individuals. Effective control is implemented through individual and collective training that ensures performance to standard. Physical controls take the form of barriers and guards or signs to warn individuals and units that a hazard exists. Additionally, special controller or oversight personnel responsible for locating specific hazards fall into this category. Avoidance controls are applied when leaders take positive action to prevent contact with an identified hazard." [798]

To summarise, there are two ways in which risk assessment techniques can be used to prioritise the recommendations from incident reports. Firstly, they can be used to rank the causes of an incident. Resources can then be deployed to focus on the generation of recommendations that address those causes with that pose the highest risk to the continued safety of an application. Secondly, recommendations might be identified for all causes without predetermining the relative importance of particular causes. Once those recommendations have been identified investigators can rank them by performing a post hoc risk analysis on the causes that are associated with those recommendations. This has the advantage of enabling investigators to increase the importance of recommendations that are perceived to address more than once cause. In our case study, Table 12.4 showed how the recommendation that 'All personnel must have confidence in their authority to immediately stop any unsafe life threatening act and exercise it accordingly' was proposed by the Army investigation to address both the Battalion S3 and the breaching team members' failure to stop the 'deviated and unpracticed' operation. Post hoc risk assessments can take this into account. This is arguably less likely if recommendations are only identified after a risk assessment has been performed on the causes of an incident.

A number of further problems complicate the use of risk assessment techniques to prioritise recommendations. The US Army [796] and Air Force [794] guidelines argue that recommendations should be associated with individual causal factors. The US Army's FM 100-14 goes on to argue that the causal factors in accidents and incidents help to identify the hazards that drive risk assessments. We have extended this argument by using these risk assessment techniques to derive priorities for the recommendations that are associated with particular causal factor. This creates problems because incidents are not, typically, the result of individual causal factors. They are, instead, the result of complex conjugations of causes. This is emphasised by the differences between the previous formula $Risk = Frequency \times Consequence$ and the more complex formulations of the partition models that were introduced in Chapter 11. The observation that incidents stem from causal complexes rather than individual causal factors has important implications for the use of risk assessment techniques to prioritise proposed interventions. If recommendation techniques focus on singular, particular causes rather than combinations of causes then investigators may fail to address systemic issues. For instance, Table 12.3 summarises several recommendations that advocate improved training as a potential remedy for the Bangalore Torpedoe incident. The combined effect of such individual recommendations might encourage investigators to consider a more systematic reappraisal of training procedures. Similarly, proposals to improve the 'safety culture' within an organisation have the potential to address many different hazards [342].

We have argued that the priority of a recommendation is determined by the risk associated with the cause or hazard that it is intended to address. This creates problems if the proposed recommendation only has a negligible effect upon a high-risk hazard. From this it follows that *the priority of a recommendation is determined by the reduction that it causes on the risk of an associated hazard or hazards.* There are, however, considerable practical difficulties involved in assessing the likely impact of a particular recommendation. This is acknowledged within the US Army guidance; "risk management is the recognition that decision making occurs under conditions of uncertainty" [798]. Uncertainty stems from several layers of subjunctive reasoning. The investigator must assess the likely probability of a hazard recurring then they must assess the likely consequences of that hazard. Finally, they must assess the potential impact that any recommendation will have on their predictions about the frequency and consequence of future incidents!

A number of important consequences stem from the notion that the priority of a recommendation can be determined by the expected reduction in the risk of a particular hazard. In particular, investigators may have to accept that the residual risk after any recommendations have been implemented remains so high that an operation or task should not be permitted to continue. For example, incident data was used to justify permanently suspending the use of the 1370-L956, flash artillery simulator, M110, during any training activity in the US Army. The 1370-L956 was "identified as contributing to numerous serious injuries of our military members during training activities and was permanently suspended from future use with units directed to turn in all unused assets" [817]. As might be expected, the overall residual risk associated with a system is determined by the maximum risk associated with a particular hazard and not the average of those risks:

> "If one hazard has high risk, the overall residual risk of the mission is high, no matter how many moderate or low risk hazards are present... The commander must compare and balance the risk against mission expectations. He alone decides if controls are sufficient and acceptable and whether to accept the resulting residual risk. If he determines the risk level is too high, he directs the development of additional controls or alternate controls..." [798].

Previous paragraphs have introduced the US Army's five stage process of risk analysis: identify hazards; assess hazards; develop controls and make risk decision; implement controls; supervise and evaluate. Previous paragraphs have described how the first three stages can be used to prioritise recommendations in terms of the difference between an initial risk assessment and the residual risk associated with both the particular causes of an incident and the more general hazards that an incident helps to identify. Of course, the residual risk that motivates the promotion of a particular recommendation will only be achieved if the remedial actions are effectively implemented. The incidents that have been described in previous chapters of this book provide some idea of how difficult it can be to ensure such conformance.

The problems of implementing recommendations can be exacerbated by the organisational and institutional boundaries that exist between investigatory and regulatory authorities. As mentioned in Chapter 4, these distinctions help to preserve the investigators' independence from those who are partly responsible for promoting an industry. One consequence of this is that the powers to ensure compliance, typically, rest with the regulators rather than the investigatory agencies. There have been notable instances in which this has resulted in recommendations not being policed or enforced in the aftermath of previous incidents [193]. Such situations have been rectified by creating a clear distinction between the roles of economic regulation and the policing of safety requirements. The follow section builds on this analysis by investigating the processes that support the implementation of particular recommendations.

12.3 Process Issues

The previous section investigated a number of recommendation techniques including the 'perfectability' approach, high level heuristics, navigation techniques including enumerations and recommendation matrices, generic accident prevention or barrier models and risk assessment techniques. These approaches are intended to help investigators identify interventions that will either mitigate the consequences of failure or will reduce the likelihood of similar incidents occurring in the future. It is important to recognise, however, that such a list of recommendations is not an end in itself. They must be validated and then presented to regulatory bodies and safety managers. They may challenge the utility of particular recommendations. The following paragraphs, therefore, analyse these additional stages that must be passed before a proposed intervention is adopted and then implemented.

12.3.1 Documentation

It is important that investigators document the recommendations that are intended to address potential problems in existing systems. This is essential if others are to implement any proposed interventions. This does not simply involve drafting guidelines to describe the proposed recommendation. In most reporting systems, investigators must also document the reasons that motivate particular findings. This is important if regulators, safety managers and other personnel are to understand the motivation for intervening in existing working practices. It is possible to identify a range of additional information that must be provided to support particular recommendations:

- *what causes or hazards does the recommendation address?*
  The opening sections of this chapter cited army and air force guidelines which require that recommendations are closely tied to particular causal factors. This is intended to ensure that as much as possible is learned from an incident; every cause should be addressed by at least one recommendation. Later sections have extended this argument by identifying recommendations, such as improvements in 'safety culture' or in training practices, that may address many different causes of a particular incident. Finally, it has been argued that incident investigations can uncover potential hazards that were not involved in a previous incident but which have the potential to jeopardise future safety. It is important for each of these cases that investigators explicitly identify the hazard that a recommendation is intended to address. Without such information it will be difficult for others to assess whether or not a proposed intervention provides sufficient protection against future failures.

- *what is the significance of the cause or hazard that a recommendation addresses?*
  As we shall see, recommendations are often passed to regulators or safety managers who must then guide the allocation of finite resources to ensure that they are implemented. From this it follows that investigators must help others to determine how to maximise their use of these resources. Risk assessment techniques have been proposed as a potential means of assessing the importance of a recommendation [798]. This can be derived from the risk associated with the hazard that a proposed intervention is intended to address. Unfortunately, a number of

problems complicate the application of this technique in ‘real world’ systems. In consequence, a great deal of subjective judgement, of skill and expertise is required in order to assess the significance of a particular recommendation. Unless such judgements are documented, however, there are few guarantees that resources will not be diverted towards relatively trivial changes whilst more significant recommendations are neglected.

- *what are the intended consequences of the recommendation?*
  Ideally, we would like to document measures that can determine whether or not a recommendation has been successfully implemented. This is easier with some recommendations than others. For instance, it is relatively straightforward to initiate plant inspections as a means of determining whether or not process components have been replaced. It can be more difficult for investigators to schedule inspections that might be necessary to determine whether a particular change has been made in a training regime. This often involves complex scheduling of site visits that can alert operators to a forthcoming inspection. There are further problems. It is generally much easier to determine whether or not a change has been made in an application process. It can be far more difficult to demonstrate that any change has had an anticipated impact upon the overall safety of a system. As we have seen, poor submission rates and reporting bias can prevent reliable conclusions being drawn from raw incident data. Investigators should, therefore, consider how to demonstrate the effectiveness of any funds that are invested in the implementation of particular recommendations.

- *who will implement and monitor each recommendation?*
  The US Army and Air Force heuristics urge investigators to identify the individuals or groups who are responsible for implementing particular recommendations [794, 796]. Investigators must not to specify *how* to implement a recommendation. This is important because investigators may lack the local expertise that is necessary to determine how best to implement a particular improvement. Similarly, the design and coordination of any changes might take far longer than the period of time that can be devoted to a particular investigation. Instead, incident reports must document *what* a recommendation is intended to achieve and *why* that objective is important. It is clear important, however, to determine *who* is responsible for implementing any proposed intervention. This individual must determine *how* to realise a recommendation from the investigators’ description of *what* a recommendation must achieve and *why* it must achieve it. If they confuse the investigators’ intentions or if they lack the resources to implement necessary changes then there is a danger that past failures will recur as future incidents.

- *establish the time-frame for any recommendation*
  The implementation of recommendations can be delayed by resource limitations, lack of managerial guidance, deliberate obstruction and so on. Ultimately, this can leave any system exposed to repeat failures if proposed changes are not introduced in time. In consequence, it is important that investigators specify *when* a recommendation should be implemented. There is a danger that this maximum time period will be seen as a target and not as an upper boundary for any remedial actions. Many investigators, therefore, provide detailed guidance on the phased introduction of particular recommendations. It is also important to monitor the implementation of key changes beyond the immediate aftermath of an incident. If this is not done then there is a danger that organisations will gradually forget previous lessons. In consequence, it is also important to consider how the monitoring of a particular recommendation might be incorporated into more routine activities.

The US Army’s Accident Investigation Handbook illustrates the way in which organisations can provide detailed guidance on the approved format for the presentation of recommendations [803]. This handbook separates its advice into three causal categories: human error; material failure or malfunction and environmental factors. There are small differences in the information that is to be recorded for recommendations that address hazards in each of these different sections. For example, the handbook requires that investigators document a range of information describing human ‘errors’. This includes a single sentence about what happened. This is then followed by a brief description of the context in which the incident occurred, for example "while conducting night convoy operations using blackout drive lights". Investigators must also identify the individual involved in the 'error' by describing their duty position, such as the OH-58D pilot-in-command or the driver of the M998, High Mobility Multipurpose Wheel Vehicle. As can be seen, such a requirement affords a degree of anonymity. Investigators must then identify the task error of omission or commission that motivates particular recommendations. These are classified according to Army standards. In particular, the accident investigation handbook recommends the error codes that are presented in PAM 385-40 [796]. These codes were used in recommendation matrices, such as Tables 12.8, 12.9, 12.10 and 12.11, that were presented earlier in this chapter. The example cited in the accident handbook is that the operator "exceeded the posted speed limit of 40 MPH by attempting to drive at 60 MPH in violation of Camp Swampy Reg 190-5 (Code 40)" [803]. This discussion of what happened then motivates an explanation of the consequences of the error. It may directly or indirectly result in damaged equipment or injury. For example, a road traffic accident may involve substantial damage to a vehicle and its driver. It can also involve injury to third parties, such as pedestrians and other drivers, as well as damage to other vehicles or objects in the vicinity of the incident. After having described the context in which an error occurred and having explained the consequences of that failure, investigators must document the reasons why it happened. In other words, they must record the findings of any causal analysis. As before, these causes must refer to the predefined lists that are provided in PAM 385-40 [796]. These are supported by a free-text description of the reasons why an error occurred: "the driver's actions were a result of a lack of self-discipline and improper supervision by the senior occupant... the driver had a history of speeding [803].

The documented 'causes' of an error help to motivate the subsequent section of the report that details the particular recommendations which are made in the aftermath of the incident. These are intended to answer the question, 'What to do about it?'. Previous sections have already described how the US Army relies upon an enumerate list of recommendations that are published in PAM 385-40. The Accident Investigation Handbook, therefore, suggests that investigators consult this document before drafting their recommendations. It is important to note, however, that recommendations should not be addressed at the task error itself but at the system deficiencies that led to the error. This approach is advocated in the handbook and explicitly encouraged in PAM 385-40 by including relatively few recommendation codes that might support a 'perfectability' approach. Recommendations must be addressed to unit level (company, troop, battalion), higher level (brigade, division, corps) or to Army level. The following format is recommended:

"RECOMMENDATION (1, 2, 3, etc.):

- a. Unit Level Action: Commander, ____ (unit): Brief all unit personnel on the facts and circumstances surrounding this accident. Emphasis should be placed on how human limitations combined with less than optimum systems and high task loading allow situations that contribute to undetected hover drifts.
- b. Higher Level Action: None.
- c. Army Level Action:
  - (1) Commander, U.S. Army Training and Doctrine Command:
    - (a) Validate requirements for automatic hover systems for all aircraft to assist in reducing task overloading.
    - (b) Validate OH-58D crew coordination requirements, especially in Tasks 10 67, 1114, 1140, 1147, and 1148 in TC 1-209, to ensure safe compliance with the requirement for both crew-members to simultaneously direct their attention inside the aircraft, especially in aircraft without automatic hover systems.
    - (c) Validate requirements for night vision systems with greater fields of view and resolution.
    - (d) Increase, within the flight-training program, emphasis on situational awareness and spatial disorientation.

  - (2) Program Executive Officer, Aviation, field upgrades to OH-58D aircraft which allow the use of the hover bob-up mode symbology in the LCD unit, even with weapons displayed in the LCD unit, and allow for adjusting the ODA intensity during low light ambient conditions.
  - (3) Commander, U.S. Army Safety Center, disseminate/publish the facts and circumstances surrounding this accident as appropriate." [803]

As mentioned, the US Army guidelines provide similar advice on how to document recommendations for other categories of failure, including equipment problems and environmental issues. In the case of material failures or malfunctions, investigators must explain what happened in a similar fashion to that described for human error. Such failures are defined to occur when a piece of equipment "did not operate as intended or designed which contributed or caused the incident". Investigators are encouraged to search for human errors or mistakes, such as a failure to follow Army standards/procedures, design criteria or manufacturing process, that may have caused the material failure. As before, it is important to document the results of any causal analysis. This is again used to identify appropriate recommendations using PAM 385-40.

Environmental recommendations follow a similar format. They are presented at the end of an analysis of the failure that describes what happened and why it happened in the manner that it did. The US Army guidelines also suggest that investigators can determine if an environmental factor should be assessed by asking 'did this factor adversely influence human and/or equipment performance; was the environmental element unknown or unavoidable at the time of the accident/injury/occupational illness?'. The explanation of why an environmental factor affected safe and successful operation often draws upon a range of disciplines. Microbursts provide an example of such a factor. They have been cited as causal factors in several recent incidents involving military aircraft. These environmental events cannot be predicted with present meteorological equipment. They are also invisible to aircraft crew-members. Such incidents show how investigators are constrained in the range of recommendations that might counter the adverse effects of many environmental factors. For example, the US Army's investigation handbook includes the following example of an Army level recommendation to deal with microburst incidents: 'Commander, U. S. Army Safety Center, disseminate/publish the facts and circumstances surrounding this accident as appropriate'. In contrast, more detailed proposals are directed at unit Commanders:

- "(a) Coordinate through the Commander, U.S. Air force, 1st Weather Group, Fort McPherson, Georgia, to establish a pro-active interface with several groups sponsoring research into the area of windshear. These groups include NASA, the National Technical Information Service (NTIS), the Federal Aviation Administration (FAA), the American Meteorological Society, the Langley Research Center, and the National Center for Atmospheric Research.

- (b) Inform all aviation personnel assigned to Fort Rucker, Alabama, that severe weather in the form of microbursts can occur from isolated thunderstorms or rainshowers and cumulus clouds that give the impression of simple rainshower clouds." [803]

A final section of the guidelines focus on the documentation of recommendations that address non-causal factors. The US Army handbook focuses narrowly on "findings that did not cause or contribute to the cause of the accident but contributed to the severity of injury or accident". An example would include a drivers failure to wear a seatbelt. This would not have caused a collision but would have significantly affected the injuried that the soldier sustained should a collision occur. This narrow definition of non-contributory factors might, however, be revised following the arguments that have been made in previous sections. For instance, non-causal factors should be extended to include hazards that have been detected during the previous analysis but that did not contribute to the particular incident under investigation. The Army handbook recommends that these non-contributory factors should each be recorded in a single paragraph; 'they are recorded to inform the command of problems that, if not corrected, could adversely affect the safety of future operations'. Recommendations that address these potential hazards are documented after recommendations that deal with human 'errors', material failures and environmental factors.

This section has argued that investigatory organisations must publish guidelines that support the documentation of particular recommendations. It is important to identify those hazards or causes that are address by particular findings. This helps to ensure that important lessons are not overlooked if potential hazards are not addressed by particular recommendations. Investigators must also document the perceived significance or importance of those hazards that are addressed by a recommendation. This information is necessary if others are to determine the best allocation of finite resources when implementing several, possibly conflicting, findings. Investigators must document the intended consequences of a recommendation. They must explain what it is intended to achieve rather than how it is intended to achieve it. This provides a degree of flexibility to engineers who must determine the best mans of implementing a particular recommendation. The documentation of recommendations must determine who is responsible for ensuring that a finding is acted upon. They should also be provided with documents that describe a potential timescale for their actions. The importance of these documentation requirements varies from organisation to organisation. For instance, in local reporting systems the investigator may also be responsible for implementing any recommendations. In such contexts, much of this information may be superfluous unless for auditing purposes. Many larger organisations, including the US Army, draft regulations and guidelines to ensure that most of this information is documented. These requirements are intended to ensure that the recipients of particular recommendations have sufficient information for them to validate any proposed changes in working practices.

12.3.2 Validation

Previous sections have focussed on techniques that investigators can use to draft recommendations that avoid or mitigate future failures. Such techniques only provide a partial panacea to the problems of incident reporting. A number of additional issues must be addressed before particular recommendations can be introduced to support the operation of safety-critical systems. For instance, there is a danger that valuable resources will be allocated to ineffective remedies. Some recommendations have been motivated by organisational politics and managerial ambition rather than a concern to address the causes of previous failures. There is also a danger that by addressing one set of problems, recommendations will inadvertently introduce other potential problems into an application. It is, therefore, important that recommendations are validated before they are implemented.

The way in which recommendations are validated can differ greatly between reporting systems. Many local systems rely upon informal meetings between the colleagues who are responsible for running the system. Large-scale systems often validate recommendations at several different levels within an organisation. Investigators may pass on the initial findings to their immediate superiors. They perform an initial check and then pass a revised version of the recommendations to their superiors and so on. Some incident reporting systems also encourage dialogues between investigatory bodies, regulatory organisations and system management. These joint meetings help to ensure that each party understands the implications of a particular recommendation. Chapter 9 has described how these dialogues can, occasionally, introduce unacceptable delays into the implementation of important safety measures, such as Excess Flow Valves into gas service lines [588].

The US Army's Accident Investigation and Reporting Procedures Handbook contains detailed guidance on the different review procedures that are to be implemented at different levels within the command structure [806]. Reports about high-consequence incidents are validated at a local review, by installation level safety-managers, by an approving authority appointed to represent the Major Army Commands and by the US Army Safety Centre. The initial review is normally conducted by the commander of the unit or by the commander of the supervisor directly responsible for the operation involved in the incident. Their must review the report and provide written feedback about whether or not they concur with the findings and the recommendations. They must ensure that any evidential data is circulated within the unit so that it can be used to inform future decision making. They are also responsible for ensuring that any immediate actions are implemented as a local level. The local reviewing officer then hands the report through the designated chain of command to the 'approving authority', see below.

There is a danger that incidents and accidents may form part of a wider pattern within a partcular installation. Similarly, there is a danger that particular recommendations that are intended to protect the operation of particular processes will have knock-on effects for the safety of other workers elsewhere in an installation. The installation-level safety manager's review is intended to identify any of these issues. The US Army reporting froms (DA 2397-R-series form, DA Form 2397-AB-R, DA Form 285, or DA Form 285-AB-R) contain special sections that are intended to help safety-managers identify these potential problems. Safety managers must review the data in these sections, not so much to validate particular recommendations, but to ensure that as much as possible can be learned from an incident. If primary and secondary investigators have missed previous incidents or patterns of systemic failure then this stage of validation is intended to identify them.

The 'approving authority' provides a further level of review within the US Army procedures. Major Army Commands appoint these representatives to accept or reject each finding and recommendation made by an investigation board. This takes place after the reports have been amended by local reviewing officials, using the procedures described above. In addition, the Safety Office of the Major Army Command ensures that the report is complete with respect to the Army guidelines [806]. Major Army Commands-level recommendations will be tracked using a computerised tracking system. At this stage, the approving authority will also be concerned to identify any additional recommendations that might be made to 'higher headquarters'. Finally, the US Army Safety Centre reviews all reports to ensure that they conform to regulatory and technical requirements. They are also responsible for maintaining the automated tracking system that the Major Army Commands use to track the implementation of particular recommendations. The Safety Centre is also responsible for disseminating information about the implementation of accepted recommendations to the relevant elements within the Army command structure.

It is important to emphasise that such elaborate validation procedures can create a number of potential problems. In particular, the responsibility for validating and implementing particular recommendations can become lost between the various exchanges that take place at different levels within the command structure. The opportunity for administrative delays is, therefore, acknowledged by guidelies that are intended to keep investigators and contributors notified about the course of the validation and implementation process:

> "Acknowledgements: upon receipt of written notification of recommendations, the responsible Department of the Army-level organisation will provide an initial response to the US Army Safety Centre within 60 calendar days as to corrective action(s) initiated or planned. Interim and follow-up reports are required every 90 days after initial response until the action(s) is closed.
>
> Return non-concurrence or rebuttals: all Department of the Army-level recommendations not accepted or implemented by the responsible command, organisation, agency, or activity will be returned to the Commander, US Army Safety Centre, with support rationale within 60 calendar days after initial notification." [806]

Local reporting systems provide a strong contrast to the elaborate procedures and mechanisms that are exploited by large organisations such as the US Amry. Peer review is often the only form of validation that is used to assess potential recommendations. These are often ad hoc, undocumented and informal. For example, many hospital-based systems hold monthly meetings between clinical and nursing staff. These discussions are, typically, unminuted. They are focussed to ensure the rapid implementation of changes providing there is general agreement about the utility of a particular proposal. There are, however, increasing pressures for such local initiatives to follow more documented processes [633, 453]. The importance of clinical audit within the medical domain and the wider public concern over high-profile accidents has led to a requirement the individuals and organisations explain *why* particular recommendations are not implemented. In consequence, the following paragraphs concentrate on the more formal mechanisms that have been exploited by large-scale systems. These may, of course, have to be scaled down to meet the more constrained budgets and scope of local systems.

Both ad hoc and more formal validation procedures must determine whether or not to accept particular recommendations. If a proposal is accepted then the review panel implicitly accepts a degree of responsibility for the proposed intervention. It is, therefore, important that they agree both with the form and the purpose of a recommendation. In consequence, many review bodies have introduced further distinctions beyond a simple accept or reject decision based on the recommendation that they have been asked to review. For example, the following quotation is part of a letter from the Commander in Chief of the US Army's Central Command. This letter reviews the recommendations that were made in the aftermath of an incident on a firing range in Kuwait. Rather than simply accepting the recommendations outright, the review approves of the intention behind the proposal but modifies it and also clarifies that the modification should not bias the implementation of the recommendation:

> "d. Recommendation 1403 provides, That appropriate administrative action be taken against the Ground Forward Air Controller. The recommendation is modified, as follows; That administrative or disciplinary action, as appropriate, be considered with regard to the Ground Forward Air Controller. The recommendation, as modified, is approved. My modification does not in any way reflect my view as to what action may or may not be appropriate. It is intended to assure the appropriate Service official of his or her complete discretion in the matter." [824]

These distinctions can be summarised as follows:

- *Accept.*
  Given the investment in time and money that is often made to support incident investigation, it might be expected that most review boards will concur with the findings of an inquiry. Unfortunately, this can be surprisingly rare. As we have seen, some guidelines explicitly argue that investigators should not consider the costs associated with the implementation of their recommendations. These considerations often prevent regulatory organisations from sanctioning the implementation of particular interventions. A host of other issues can prevent review boards from accepting the findings of incident investigators. For example, the members of these boards typically do not take part in an initial investigation. It can, therefore, be hard for them to follow the detailed causal arguments that motivate particular recommendations. Review boards, therefore, often request further clarification or additional forms of evidence before they will accept many proposed interventions.

- *Accept with provisos.*
  Most review boards do not immediately accept all of the recommendations that are proposed by investigators. Instead, they may request additional evidence to support a causal analysis. Alternatively, review boards may propose alternative causal explanations that, if proven, would support other forms of intervention. Even if a recommendation is accepted, review panels may advise that its implementation is delayed or staged. Such ammendments can be motivated by the financial constraints, mentioned above. They can also reflect the pragmatic problems of ensuring conformance to any proposed changes in working practices and equipment. These provisos are typical of reporting systems in which investigators are independent from any regulatory function. They also characterise more local systems in which investigators must secure the support of higher levels of management before any commitment can be made towards increased investment. In such circumstances, review boards can accept recommendations 'subject to approval' from upper management.

- *Reject.*
  Review boards, typically, exploit one of several standard 'forms' of argument when attacking investigators' recommendations. The first line of attack rejects the arguments that investigators make during the causal analysis of an incident. For example, review boards can use variants of the counterfactual arguments proposed during a causal analysis by suggesting that an accident would still have occurred even if particular recommendations were implemented. Alternatively, it might be argued that proposed interventions only address the specific causes of an incident but fail to address more general failures. A second line of attack can be based around the risk assessment techniques that were introduced in previous paragraphs. It can be argued that the expected frequency or consequences of any future incident would be too low to justify the expenditure that is required to implement the investigators' recommendations.

- *Reject with provisos.*
  Review boards must exercise a considerable degree of caution when rejecting the recommendations in an incident report. They run the risk of alienating the investigators who constructed such documents. There are obvious dangers in praising a review board for their careful use of resources if an incident does not recur within a given time period. Such rejections can also create a form of implicit responsibility should an incident recur. If an incident does recur then it can be argued that the failure might have been avoided if they had only approved the proposed intervention. It is, therefore, particularly important that review bodies document their reasons for rejecting a recommendation. In practice, this often leads to partial rejections or a refusal to implement a particular finding until some other condition is satisfied. This condition may involve eliciting additional evidence. It might also involve a commitment to perform additional studies should further incidents be reported.

- *Referral.*
  Given the potential consequences of rejecting a recommendation and the possible costs associated with implementing some proposed interventions, it is hardly surprising that many review boards defer to another authority rather than reach a premature decision. Often validation exercises result in panels deciding that they are not competent to reach particular decisions. Alternatively, they may accept the high-level arguments associated with a particular recommendation but refer to another body who must then develop a more detailed implementation plan. This is an interesting strategy because that body then assumes partial responsibility should the costs exceed expectations or the implemented remedy fail to prevent future incidents.

This list illustrates the range of outcomes that validating bodies might consider when assessing a recommendation. It is remarkably rare for a review panel to accept every recommendation without some caveat or proviso. Most validation exercises accept some proposals, reject a few recommendations and request that the remaining proposals be amended in some form. It is important to note, however, that a number of comments can be made about these general remarks. For example, many incident and accident reporting systems exploit a hierarchical validation process where review committees at a lower level in an organisation review the investigators' proposals before they are validated at a higher level. At each stage in this validation process it becomes less and less likely that higher authorities will reject a recommendation that has been accepted at a lower level. A cynical interpretation of this process might be that political and organisational pressures can help to mould recommendations into an acceptable format before they are presented to the highest levels within an organisation. A more favourable view is that upper management are less likely to question the detailed operational decisions of their subordinates.

A recent incident involving an Australian Army cadet helps to illustrate how different individuals and groups play different roles in the validation of particular recommendations. This incident occurred when a regional cadet unit were completing an exercise in which they had to swim to retrieve an object from a boat that was some twenty meters from the shoreline of a Dam. The cadets were wearing their army fatigues and boots. Several of them became entangled in weed beneath the surface of the water. One cadet became exhausted and went under the water approximately seven meters from the shorelines. Efforts to rescue him were unsuccessful. Arguably the highest level of validation for the Board of Inquiry's findings came from the Hon. Bruce Scott MP, Australian Minister for Veterans Affairs and from Dr Brendan Nelson, Parliamentary Secretary to the Minister for Defence . They concluded that the Board "conducted a thorough and open investigation into the circumstances" surrounding the incident [731]. They agreed with the Boards finding that the "swimming activity was not authorised by (the) Army and that there was inadequate supervision or monitoring of the Army Cadet Corps activity". In consequence, they took actions to suspend all swimming activities conducted in areas other than supervised swimming pools were immediately suspended, and will continue to be so, until a new policy on swimming activities is issued. They also implemented a review of the Australian Services' Cadet Scheme policy on safety, risk analysis and activity clearance by the Defence Safety Management Agency.

Such actions illustrate the way in which a final stage of validation is usually performed by organisations that exercise budgetary or political control over the implementation of particular recommendations. Their approval is required in order to approve the investment that may be required to support large-scale change. They must also provide the political support that is often necessary to implement what are often unpopular 'systemic' changes to establish working practices [701]. It is important to note, however, that such press statements and ministerial announcements represent the final stage in a range of more detailed validation activities. For example, the Australian Army's Board of Inquiry into the previous incident initially presented its findings to the Chief of Staff, Headquarters Training Command. He then issued a detailed appraisal of their findings. These illustrate the different forms of response that were sketched in previous paragraphs. For example, some of the Boards findings were accepted without comment:

> "I accept the Board of Inquiry finding that Cadet Sperling drowned as a result of a combination of factors namely, the amount of weed in the water, the depth of water, the wearing of GP boots (with socks) and Disruptive Pattern Camouflage Uniform (DPCU) clothing whilst in the water and the absence of safety devices (such as flotation vests) and inadequate safety precautions for the swimming activity. These factors contributed to Cadet Sperling's drowning. The wearing of GP boots and DPCUs whilst swimming or treading water is a difficult activity for persons of average physical fitness. A swimming activity undertaken by cadets as young as 13 years with unknown fitness levels and unknown medical conditions in the circumstances existing on 18 Nov 00 at the Bjelke Peterson Dam, was inherently dangerous." [33]

This acceptance illustrates the way in which validating bodies do not simply consider the recommendations that are issued by investigators. Review boards, typically, begin by assessing the evidence, the course of events and the causal analysis that are presented in the opening sections of most reports. For example, the Chief of Staff disagreed with the Board's analysis of one of the causal factors that was cited as a contributory factor in the incident:

> "I do not accept the finding of the Board of Inquiry that Corporal (Army Cadet Corps) ____ was not fully qualified as an instructor of cadets in the Army Cadet Corps in accordance with the Army Cadet Corps Policy Manual. Corporal (Army Cadet Corps) ____ had completed the Instructor of Cadets Course and First Aid Course in compliance with the Army Cadet Corps Policy Manual and was qualified as an Instructor of Cadets." [33]

Such validation actions illustrate the importance of explicitly documenting the causal findings that support particular recommendations. Without such analysis, it can be difficult to determine which recommendations might be affected by the review board's rebuttal of the investigators' analysis. It is for this reason that Tables 12.6, 12.8, 12.9, 12.10 and 12.11 were introduced to provide a bridge between the products of a causal analysis and the interventions that are intended to safeguard future operation. Such documentation can help investigators to determine whether or not a recommendation must be abandoned after such a rebuttal. If, for example, a recommendation is supported by several lines of causal analysis then it may still be retained even though one line of argument has been challenged.

If reviewers accept that incident investigators have identified a cause of the incident then they may continue their validation by asking whether or not that cause is 'adequately' addressed by the proposed recommendation. At first sight, this might seem to be a relatively trivial task that should be based around an engineering assessment of whether or not an incident is likely to recur if a recommendation is implemented. As we have seen, however, such subjunctive reasoning is fraught with problems. Many of these relate to the psychological processes involved in reasoning about alternative possible futures without the support of some underlying model of formal reasoning [401]. Other problems stem from the way in which some recommendations are not intended to entirely avoid future incidents but to control or mitigate their consequences. The effectiveness of these measures often depends upon the nature of any future incident and this, in turn, may depend upon other defences functioning in the manner intended. As we have seen, however, many incidents stem from the failure of these 'defences in depth' [701]. Further problems arise when recommendations have social or political consequences that can prevent review bodies from adopting them. For example, the Chief of Staff, Headquarters Training Command could not accept one of the recommendations that would have had considerable implications on the size of the Australian Army's Cadet force: "I do not accept the Board of Inquiry recommendation (Reference A para 268(d)) that cadets suffering from asthma should be required to comply with Army recruiting standards" [33]. Such a recommendation would reduce the likelihood of future incidents. It would also sacrifice some of the wider objectives that motivate the Army and the Department of Defence to run the Cadet Force.

Previous sections have explained why it can be relatively rare for validating bodies to accept the recommendations of incident investigators without raising caveats and objections. There are, however, examples of proposed interventions that are accepted in this way. It is important that the review board explicitly documents the extent of their agreement so that there can be no subsequent disagreement about what was intended by their approval for particular measures. For example, the following review paraphrases the Board of Inquiries recommendation and uses their paragraph reference scheme, Reference A para 268(f), to make sure that the reader can trace their agreement back to the original proposal:

> "I accept the Board of Inquiry recommendation (Reference A para 268(f)) that the Application for Activity Approval be forwarded through the cadet unit's foster unit with the provision for comment and then on to the respective Regional Training Center for consideration. On approval or rejection, a copy of the Activity Approval Form should be returned via the foster unit who should then confirm the availability of requested equipment and other support. The revised arrangements are to be incorporated into the Army Cadet Corps Policy Manual. Action: COMD Army Cadet Corps by 14 Mar 01." [33]

Previous paragraphs have described how review bodies can respond in several different ways to the recommendations that are proposed in incident reports. They may accept them, reject them or request modifications. They may also defer comment and request additional evidence or support from others at different levels within an organisation. The following list uses the previous analysis to derive a list of requirements that might guide the validation of recommendations in incident reporting systems:

1. *Clearly identify each stage of the review process.*
   There are increasing pressures, especially within certain sectors of the Healthcare and transportation industries, to ensure that recommendations are not dismissed without due consideration. One consequence of this is that any proposals must be subjected to a clear and coherent review process if they are not to be implemented. From this it follows that each party in an investigation must understand the nature and extent of each validation. In particular, it is important that time limits be associated with each stage of a review so that investigators, regulators and contributors can track the progress of a report towards implementation.

2. *Establish that the report is complete.*
   Given that many incident reporting systems cover diverse geographical and functional areas, it is likely that some reports may omit important details about an incident. If such reports are dismissed late in the review process then there is a danger that important insights will be ignored. It is, therefore, important that an initial validation ensures that any potential report is considered complete so that any consequent recommendations will not be immediately dismissed. For instance, checks may be conducted to ensure that all relevant evidence is available and is cited correctly. Other forms of integrity check can also be carried out. For instance, if the US Air Force guidelines are followed then each recommendation must clearly identify an initial implementation route.

3. *Validate the evidence.*
   Review boards must ensure that evidence is cited in a consistent manner and that all of the necessary data about an incident has been presented in an incident report. There is an increasing recognition that complex incidents often stem from interactions between systems

failures, human 'error', managerial problems and so on. Less 'severe' incidents often cannot command the resources that are required to fund multi-disciplinary investigations. It can, therefore, be difficult for investigators to identify all of the information that might be relevant to an incident. This is especially true when individuals are unaware of similar incidents in other units or regions. In consequence, review boards must satisfy themselves not only that the relevant information has been collected but that it is also presented in a fair and impartial manner within the body of the incident report. Chapter 14 describes some of the pragmatic problems that can arise when attempting to satisfy such an abstract requirement.

4. *Validate the causal analysis.*
   Chapters 10 and 11 have described a range of techniques that support the causal analysis of adverse incidents. These approaches provide procedures to guide the analysis of adverse occurrences. They also depend upon a range of subjective decisions that must be validated. Even within the formal systems of reasoning, investigators must identify those elements of an incident that are to be represented within the abstractions of a formal logic. It is also important to emphasise that none of these techniques is 'error proof'. The correctness of any causal reasoning must, therefore, also be verified. Any omissions or errors at this stage in the analysis can result in recommendations that fail to address the causes of an incident.

5. *Validate each recommendation.*
   This chapter has reviewed a range of heuristics that can be used to validate particular recommendations. For example, investigators may lack the time and the experience necessary to identify the best means of implementing particular recommendations. It is, therefore, important that any proposals should focus on what is to be achieved rather than the particular mechanisms that will be used. Similarly, we have argued that clear timescales must be associated with each recommendation so that their implementation is not indefinitely delayed. Proposed interventions should focus on specific actions rather than on additional studies that may or may not identify potential safeguards. It is important that review bodies consider these various heuristics when validating particular requirements. Clearly, there may be instances in which some of these guidelines cannot be satisfied. For instance, if it would be dangerous to impose additional requirements without further investigations. Validation authorities must, however, satisfy themselves that there are indeed good reasons for violating these recommendation heuristics. This analysis must also consider any priorities that are associated with any proposed interventions. The risk analysis techniques, described in previous sections, often depend upon subjective assessments both of frequency and consequence that can have a profound impact upon any subsequent resource allocation.

6. *Document the reasons for any rebuttal.*
   There can be profound implications if a review body decides not to accept a particular recommendation. If a similar incident occurs in the future then they may be blamed for opposing a necessary safety improvement. It is, therefore, essential that some auditable justification should be recorded to support such decisions. This argument applies to the rebuttal of particular recommendations. It is also important to document any challenge to the evidence and any causal analysis in an incident report. For example, if a line of analysis is questioned then it is important to ensure that any associated recommendations are not supported by alternate causal arguments. If the recommendation is dismissed without such an additional check then there is a danger that s potential cause of future incidents will not be addressed by any proposed safeguards.

7. *Validate implementation plans.*
   The next section will identify some of the problems that can frustrate the implementation of recommendations once they have been approved by validating bodies. It is important, therefore, that review organisations should consider these potential barriers when assessing particular recommendations. If they request resources that cannot be made available at a local level then the validating authorities must provide some means of ensuring that additional resources are provided. If such resources cannot be found then they must either recommend

that a proposal be redrafted or, in extreme cases, that production should be halted until some remedy is identified. This validation activity does not simply focus on the staff and equipment that may be necessary to perform any changes to an application. It also focuses on the key personnel who must supervise those changes. In particular, the implementation of particular recommendations should not impose additional burdens that may result in other forms of failure being introduced into a system.

8. *Initiate recommendation tracking.*
   The problems that exacerbate the implementation of potential recommendations have motivated many organisations to create automated tracking systems. These enable safety managers to request and review reports from individual units as they are scheduled to adopt any changes in their working practices. These tracking systems are often integrated into the final stages of validation. Once a recommendation has been approved for implementation then an entry is created in the tracking system. This is tailored to reflect the timetable and monitoring responsibilities that have been proposed by investigators and approved by successive reviews.

The validation of particular recommendations provides no guarantees that they will ever be implemented. The complexity of many safety-critical applications can provide numerous barriers to the introduction of process improvements. It can be difficult to ensure that key personnel understand what they must do in order to avoid future incidents. Similarly, it can take months and even years before obsolete components are removed from a system. Even within the best resourced systems, engineers are often found to retain stocks of spare parts that have been condemned in previous incident reports [806]. The following section, therefore, briefly considers some of the challenges that must be addressed when investigators and safety managers must implement particular recommendations.

12.3.3 Implementation

The implementation of recommendations involves the development and monitoring of a corrective action plan [571]. These plans are prepared by individuals who are, typically, appointed by the most senior validation board. These 'implementation officers' may or may not have been involved in the initial incident investigation. Their action plan must explain how they propose to address all of the recommendations that have been accepted following ammendment and clarification. Each item in the action plan must address the following questions:

- *What causes are addressed?*
  In order for managers and operators to understand the importance of a corrective action, information should be included about those causes of previous incidents that are to be addressed by a particular intervention. NASA explicitly recommend that portions of a recommendation matrix should be included with an action plan [571]. This may, however, prove to be too cumbersome a requirement for smaller scale systems.

- *What is to be done?*
  The recommendations that are validated by review boards should describe what is to be achieved without describing how any particular requirement will be satisfied. Hence, this information can be directly derived from the final version of a recommendation that is approved by any review board.

- *How is it to be done*?
  It is important that managers and operators can plan how to satisfy a particular recommendation. As mentioned above, this detailed information need not form part of the documented proposal that is validated by review boards. It must, however, be documented in an action plan that can be approved prior to implementation.

- *Who is responsible?*
  The proposed action plan must clearly identify who is to implement any intervention. This can involve a detailed consideration of which branch of an organisation or subcontractor is responsible for ensuring that a corrective action is completed.

- *What are the wider consequences of any corrective action*?
  The corrective action plan must consider any wider implications that result from the implementation of a particular recommendation. Previous sections have mentioned how some interventions can increase the risk of other forms of incident. Such trade-offs may have to be accepted if the benefits of preventing other forms of failure are perceived to outweigh this collateral risk. Corrective action plans must also review any wider process changes that may be necessary following the implementation of a recommendation.

- *How will the corrective actions be tracked?*
  It is important to ensure that corrective actions are implemented correctly if they are to have the intended impact upon overall system safety. An action plan must, therefore, consider how any interventions will be tracked. This analysis should ideally provide for interim status reports and for documentation to confirm the completion and closure of corrective actions.

The Canadian Forces provide an example of such action plans being used to direct the implementation of particular recommendations, known as needs assessments [149]. They encourage the development of specific implementation programmes that are intended to meet these needs assessments. In addition to the high-level requirements mentioned in th previous list there is also a concern to ensure that any action plan considers an appropriate range of potential implementation mechanisms. Implementation offers 'border on the negligent' if they only propose solutions that involve additional training. Improved tools, procedures and job-aids provide alternative solutions to inadequate knowledge or skills. It is ironic, however, that if these planned changes are not implemented then there may be future incidents are likely to be reported as training failures.

As with the approval process that is used to validate individual recommendations, implementation officers must identify a timetable both for the drafting and the approval of an action plan. For example, it might be specified that these actions should be completed within 30 working days of a validation panel accepting a particular recommendation unless they provide a written justification for extending the deadline. As mentioned, implementation plans are often not developed by investigators. It is important. however, that any action plan should be passed to them so that they can provide high-level feedback about whether or not the proposed intervention will fulfill their particular recommendations. Copies of an action plan may also be passed by the validating panel to safety managers and to regulators for further review. Their comments must be considered by the validation panel within the timescales, described above.

If an implementation plan is rejected by the validation panel then it is returned to the responsible organisation for revision and resubmission. As before, a timescale for resubmission must be developed to ensure that potential safety improvements are introduced as soon as possible, It is important to emphasise that this process of working out how to implement a particular recommendation can help to uncover further recommendations that might not have been considered during an initial investigation. For example, 'cook-off' incidents occur when the heat that is generated by a gun can cause premature firing of ammunition. A series of incidents persuaded the US Army to focus on the M60 machine gun. During a more detailed analysis of potential solutions to this problem, it was realised that 'cook off' incidents also affect a range of other weapons that had not been considered during the initial analysis [820].

If a plan is accepted then the implementation officer must initiate the proposed corrective actions, for instance by putting out any proposed work to tender or by disseminating relevant safety information. In larger organisations, these actions will, typically, be performed in close collaboration with safety management. In smaller organisations, an action plan may simply be approved by higher management and then be initiated by the staff running the reporting system. In either case, audit actions are often introduced so that review bodies can determine whether corrective actions have been implemented and whether they can be shown to produce the desired effects. Previous sections have mentioned the difficulties of measuring safety improvements when adverse incidents are likely to be rare events. A range of further problems complicate these audit activities. For examples, the individuals and groups who are responsible for executing an action plan may discover that certain actions are unnecessary or unwise. In such circumstances, the implementation officer must seek approval to alter the implementation plan. Such changes must be well-documented and validated by a review board and by safety management before they can be accepted.

It is important to determine who is responsible for monitoring compliance with particular safety recommendations. In smaller-scale systems, this is likely to be the same person who is responsible for ensure the implementation of any corrective actions. In larger scale systems, this monitoring function is more likely to be performed by an independent safety manager who must report any concerns about non-compliance to the validating panel. This feedback is necessary for several different reasons. For example, it can be difficult for validating bodies to identify whether or not a proposed intervention will be effective unless they are informed about the success or failure of previous initiatives. Similarly, review bodies may be able to act if they identify patterns of non-compliance within particular geographical areas or functional units. Chapter 15 will discuss the problems of interpreting and acting on such feedback in greater detail.

The implementation officer uses the responses from any monitoring together with any independent analysis from safety managers to determine whether or not it is possible to close a corrective action. Some organisations require approval from the validation or review body [571]. This approval can be obtained once the implementation officer submits a final incident review. This review includes the investigators' incident report, the corrective action implementation plan and a list of any additional lessons that have been learned from an adverse occurrence. The review should also document any significant departures from the approved implementation plan as well as any non-compliance concerns that had to be addressed. Final review documents should be archived for future reference. This is increasingly done using electronic databases and information retrieval systems. Such tools enable investigators and safety managers to automate the search tasks that can be used to identify previous recommendations for similar incidents. Chapter 14 considers a range of potential technologies that can be used to support these tasks.

The US Air Force provide a specific example of the generic final review document mentioned in the previous paragraph. Their Air Force Instruction (AF 91-204) sets mandatory standards for incident and accident reporting [794]. This refers to a memorandum of final evaluation. The Headquarters Safety Centre must draft one of these documents for each high-criticality incident that is reported to them. This is an important caveat, clearly the extensive implementation procedures mentioned in previous paragraphs might place too high a burden on organisations responding to low criticality events. If individual operators and investigators felt that the procedural burdens outweighed the potential benefits from a particular recommendation then their might be a tendency to suppress or limit the number of proposed interventions. The Air Force, therefore, is careful to specify when these procedures must be followed. For instance, a Memorandum of Final Evaluation must be prepared for Class A and B incident reports even when the requirement to produce a formal report has been waived. Class A mishaps include 'failures' that incur cost of $1 million or more. This classification covers fatalities, permanent injuries or the loss of an aircraft. Class B mishaps include 'failures' costing between $200,000 and $1 million. Events may result in permanent partial disability or hospitalisation.

The Memorandum of Final Evaluation collates input from various sources including the Major Commands that convene an investigation, the commander of the mishap wing, statements from individuals and groups who are cited in an incident final report and so on. It is intended to provide an overall assessment both of the incident report and of any subsequent responses to the investigators' findings. The US Air Force procedures also state that the Headquarters Chief of Safety must publish these memorandum using an electronic database (AUTODIN) and the Defence Messaging System. At this point, the memoranda become the "official Air Force position on findings, causes and recommendations" that relate to the incident [794]. The Headquarters Chief of Safety, therefore, explicitly validates the recommendations that are embodied within the memorandum through this act of publication via these information systems. Any associated actions become active and must be executed by the named agencies that are associated with each recommendation. Suspense dates are also associated with these actions. Action agencies must report on completed actions or on progress toward completed actions by that date.

All agencies and organisations within the Air Force are required to review each Memorandum of Final Evaluation to determine whether any of the deficiencies leading to the mishap apply to their commands. This involves a filtering process in which each memoranda is forwarded by a receiving officer to the technical units that might be affected by any particular recommendation that is contained within it. The directors of these units review the memoranda to determine whether or not they are applicable to their systems an working practices. If they are then changes are initiated at this level. The incident reporting process does not finish with the local implementation of any recommendations in a Memorandum of Final Evaluation. Mishap Review Panels must be established within individual commands to ensure that recommendations continue to be addressed. The regulations require that these panel meet *at least* once every six months. These meetings are intended to ensure that preventive actions are implemented and that all parties review the status of open recommendations. Recommendations must remain open until Headquarters Safety Officers agree that either all recommended changes to publications have been made and the updated versions are issued or the recommended modifications have been completed on all applicable systems or that all recommended studies and evaluations have been completed and that actions on all validated requirements have been closed. It is possible for recommendations to be closed if they are considered to be impracticable within existing operational constraints or cost parameters. Similarly, a recommendation can also be closed if an item is removed from service. Such actions must again be validated at a central level so that the outcome is recorded in the electronic information systems, mentioned in previous paragraphs.

As mentioned above, these various reporting procedures apply to major incidents and accidents. A less formal approach is permitted for less serious mishaps. For example, an incident description can be drafted instead of the more formal incident report. It is important to note, however, that these descriptions must still be validated at the Major Commands level; “While (these) mishaps are not catastrophic, they are serious enough to require reporting on an individual basis and recommendations resulting from them require effective management”. These less critical mishaps are not tracked by the Memorandum of Final Evaluation process, described above. The Air Force, therefore, introduces additional requirements to ensure that lessons are learned from the analysis of these incidents. The final description, mentioned above, must outline all of the local actions that were taken after an incident. As we have seen throughout this chapter, these remedial actions must be explicitly related to the causal findings that they are intended to address. These documents must also report any actions that are planned but not yet completed. Estimated completion dates must also be provided. These reports are also intended to provide local units with an opportunity for eliciting central support should it provide necessary in order to implement a particular recommendation.

12.3.4 Tracking

This book focuses on two different levels of tracking or monitoring within incident reporting systems. The first of these activities ensures that operators and managers conform to the individual recommendations that are made in the aftermath of incidents and accidents. We refer to this as recommendation ‘tracking’. The second of these activities ensures that incident reporting systems as a whole are having their intended effect on the safety of an application process. We refer to this as the ‘monitoring’ of a reporting system. This section provides a brief overview of recommendation tracking. Chapter 15 provides a more detailed analysis of system monitoring.

Previous pages have described how implementation action plans must be developed if high-level recommendations are to protect the future safety of complex, application processes. We have also described how electronic databases and messaging systems have been used both by the US Army and Air Force to track outstanding actions plans until they are closed. Such systems provide a particular example of more general techniques that have been developed to help implementation officers track the progress towards achieving particular recommendations. These approaches must address a number of problems that can limit the effectiveness of any implementation plan. For instance, intended recipients may not receive a plan. Tracking systems must determine whether or not all appropriate personnel have access to the information that is necessary in order for them to implement a particular plan. This might seem to be a trivial requirement given the sophisticated communications infrastructure that supports many complex, organisations. As we shall see, many incidents recur because these communications systems are not completely reliable. For instance, paper-based instructions are frequently lost or destroyed. This creates particular problems when information must be passed between different shifts or teams of co-workers. Electronic information systems often suffer from usability problems that can prevent staff from accessing the information that is necessary for them to revise previous working practices. Technical problems and server loading can also prevent uses from accessing necessary information. Further problems stem from the difficulty of keeping up with the number of implementation plans that affect the many different items of equipment that particular members of staff may be responsible for. For instance, the US Army issued at least eight revision requests for the M9 Armoured Combat Earthmover manuals in a single month in 2000: TM5-2350-262-10, TM5-2350-262-10HR, LO5-2350-262-12, TM5-2350-262-20-1 & 2, TM5-2350-262-20-3, TM5-2350-262-34, TM5-2350-262-24P, TM5-2815-240-34&P [810]. These were published in paper form and disseminated via the Army Electronic products Support Bulletin Board (http://aeps.ria.army.mil/). In addition to these sources, Armoured Combat Earthmover operators also had to monitor at least two separate web sites (http://ncc.navfac.navy.mil and http://www.tacom.army.mil/dsa/) that contained further information about modifications and revised operating procedures for their vehicles. The difficulty of following all of the implementation plans and revised regulations that affect particular tasks can also be illustrated by the US Army's explosives safety policy. Between September and December 1999, the Office of the Director of Army Safety and the Office of Deputy Chief of Staff, Logistics issued revised guidance on loading Bradley Fighting Vehicles, on the Storage of Operational, Training and Ceremonial Ammunition in Arms Rooms and on Explosives Safety Site Plans for Ranges. Each of these involved major changes in the way that safety managers and operating units conducted many 'routine' tasks. For instance, the revised guidance on loading ammunition into the Bradley Fighting Vehicles gave the following explosives safety guidance:

> "If a BFV is uploaded with only 25mm ammunition and other small arms ammunition, with the hatches and ramp closed, then that BFV is considered heavy armour. The heavy armour qualification allows such a BFV to have reduced quantity distance separations. Uploading with TOW missiles or other high explosives items removes the allowed reduction in quantity distance." [819]

These revised policies and procedures were published via the the US Army's Explosives Safety Website (http://www.dac.army.mil/es/). However, the recipients of these revised guidelines were also warned that they were minimum guidelines and that even if they followed them they may also be in contravention of more restrictive practice regulations enforced by individual Major Army Commands; "before personnel act on these policies, personnel should check with their MACOM safety offices to see if MACOM policy mirrors Army policy" [819]. This duplication of authority creates considerable problems for the operators and managers of complex, safety-critical systems. This interaction between local requirements and the recommendations from central incident reporting systems complicates the problems of ensuring conformance with safety requirements. Tracking must, therefore, assess whether operational units meet the minimum recommendations proposed by an implementation plan. It must also determine whether those units meet the more stringent requirements that are often imposed when local units seek to enforce those recommendations.

A further purpose of tracking is to ensure that operators and managers receive correct information about revised operating procedures. Many reporting systems translate the recommendations that are embodied within implementation plans into more accessible formats. For example, the US Amry publishes information about such changes in its *Countermeasures* magazine. Very rarely, mistakes can enter into a recommendation as it is translated between an implementation plan and the story that is disseminated through these publications. Such errors have important safety implications if they are not detected either by feedback from the recipients of this information or through careful tracking by the operators of the reporting system:

> "Thanks to all the sharp-eyed readers who noticed that we published the incorrect maximum allowable speed for the M939A2 trucks in last month s Countermeasure . In the article The Rest of the Story on page 12, the correct sentence should read, '...the board checked the Army Electronic Product Support Bulletin Board via the Internet website http://aeps.ria.army.mil/and discovered that there are two safety messages (GPM 96-04, 131807Z and SOUM 98-07,081917Z) restricting the maximum allowable speed for

> M939A2 trucks to 40 mph (not 45 mph as previously stated) until antilock brakes and radial tires are retrofitted. We're sorry for this error." [808]

Even if the intended recipients of an implementation plan successfully receive information about revised working practices or material changes, there is no guarantee that they will act upon them. Chapter 2 has described the problems of identifying the reasons that motivate non-compliance with safety instructions. Some incidents are due to deliberate violations; operators may not understand the safety implications of a failure to comply with particular instructions. Other incidents stem from the operators' failure to understand the procedures that are required of them. These problems can be illustrated by a recent incident in which the right-side track of an M113 Armoured Personnel Carrier snapped. This prevented the driver from steering effectively. It also prevented any braking maneuvers which increased the vehicle's pull to the left. Subsequent examination showed that the pin on one block had worn through the metal parts that held it within the adjacent track block. A deep gouge in the hull and significant wear patterns on various track parts indicating a history of improper track maintenance. The M113 crew had all of the necessary tools and manuals to identify the problems. However, neither they nor the platoon leadership nor the company commander ensured the proper implementation of preventive maintenance procedures (PMCS) and revised operating regulations (DA PAM 738-750).

It is important to emphasise that such incidents often stem from multiple failures in the dissemination of safety-related information. An individual's failure to act on a particular implementation plan can have consequences that are compounded by their lack of information about other safety issues. For instance, previous incidents had also resulted in a maximum speed limit of 25 miles per hour being imposed on tracked vehicles, such as the M113, for the type of road that the crew was driving on. The driver did not know these limits, and neither the vehicle commander nor the squad leader traveling behind him took any action to make him slow down; "excessive speed contributed to the track failure and to the rate of turn of the M113, which resulted in roll-over" [804]. Such incidents are important not simply because they reveal the problems of ensuring compliance with the recommendations that have been made following previous incidents. They also illustrate particular problems in the dissemination of information about the associated implementation plans. It is, therefore, important that the managers of incident reporting systems track the analysis of future incidents in order to assess whether or not previous recommendations are being disseminated and acted upon by operational units. Previous paragraphs have described how the recipients of an implementation plan may either explicitly refuse to revise their procedures or may neglect to follow their requirements. In other situations, personnel may be motivated to comply with an implementation plan but they may lack the necessary resources to follow its provisions. Necessary resources can include the time and skills necessary to perform new procedures. They also include any new components that are identified in particular recommendations. Finally, the recipients of an implementation plan may lack the financial resources that might otherwise be used to make-up any shortfall in other resources. Ideally, such problems will have been considered and addressed during the development of an implementation plan. It would, however, be unrealistic to assume that such preparations would obviate the need to track the recipients' ability to satisfy the recommendations in these plans.

There are situations in which the tracking of particular recommendations can reveal concerns about the effectiveness of an implementation plan. During 2000-2001, the US Amry introduced an Improved Physical Fitness Uniform (IPFU). This was intended to offer improved comfort during exercise. It was also intended to reduce accidents and incidents through the incorporation of reflective material into the uniform. Many of the personnel who were issued with these uniforms were clearly motivated to conform with these joint requirements; to increase personal comfort and ensure visibility during exercise. The Safety Centre, therefore, received several enquiries about the effectiveness of the improved uniform's reflectivity. Subsequent investigations found that the uniforms met their intended specification and the comments were not triggered by either a design or production defect. In consequence, the uniforms were not recalled in response to the end-users' concerns. Instead, the Safety Centre emphasised that the uniform was not intended to be a replacement for a luminous safety vest [811]. Such incidents are instructive because they contrast strongly with the use of implementation tracking to detect violations. In this case, staff were concerned to meet the recommendations that informed the development of the improved uniforms. They felt, however, that the improved designs did not, however, offer the necessary degree of protection. The US Army safety Centre's response is also instructive. Instead of recalling the uniforms, their analysis of the end-users comments revealed additional safety concerns. Personnel were potentially relying on the protection offered by the uniform's reflectivity rather than wear a safety vest.

Implementation officers must track whether or not these validation and dissemination processes have introduced undue delays into the implementation of safety recommendations. This can be determined if a number of similar incidents occur before necessary changes are made to working practices or to process components. Such tracking activities can also reveal a converse problem in which recipients receive warnings well before they can act upon them. This occurs, for example, when advisories are issued for equipment that has not yet been received by its potential operators. In such circumstances, there may be an assumption that such warnings do not apply to their current tasks and hence they may be ignored. This can be illustrated by the findings of an incident involving one of the US Army's M939A2 wheeled vehicles on a public road [812]. Weather and road conditions were good and the vehicle obeyed the planned convoy speed of 50 miles per hour. In spite of this, the driver of an M939A2 failed to prevent the trailer that he was towing from 'fish-tailing' as he started to descend a steep hill. One of the tires on the trailer blew and the truck rolled off the road. The subsequent investigation determined that the tires were well maintained and showed no defects. Witness statements and expert testimony confirmed that the vehicle was not exceeding the approved speed limit. The investigation board's maintenance expert asked if the unit was aware of any Safety-of-Use-Messages or Ground Precautionary Messages on the vehicle. At first, unit personnel said no. They had only recently receivied their first two M939A2 trucks as replacements for older models.

> "At that point, the board checked the Army Electronic Product Support Bulletin Board via the Internet website http://aeps.ria.army.mil/ and discovered that there are two safety messages (GPM 96-04, 131807Z and SOUM Investigators Forum 98-07,081917Z) restricting the maximum allowable speed for M939A2 trucks to 45 mph until antilock brakes and radial tires are retrofitted. Further interviews with unit maintenance personnel determined that they had seen the messages when they came out. However, since the unit did not, at that time, have any M939A2 trucks, they did not inform the chain of command. The lesson here is whenever your unit receives new equipment; it is good practice to check all relevant Safety-of-Use-Messages and Ground Precautionary Messages to ensure that you and your personnel operate the equipment safely." [812]

Such incidents illustrate the problems that can arise when attempting to ensure that implementation plans continue to be followed in the aftermath of previous failures. As we have seen, many modern organisations are characterised by their ability to change in response to their environment, to market opportunities and in response to technological innovation. This has several important consequences for those who must track the implementation of safety policies. New devices will be introduced into new working contexts. Those devices may be subject to previous recommendations that must be communicated to the operators who must employ them within these new contexts. Similarly, new devices may interact with other components or working procedures that were themselves covered by existing recommendations. These changes can force revisions to existing guidelines and procedures. It is also important to stress that many organisations benefit from a dynamic workforce that moves between different production processes and regional areas. These workers carry their skill and expertise with them. There is considerable potential for them to apply procedures and regulations that were appropriate in their previous working context but which can be potentially disastrous in their new environment. In consequence, safety managers must typically find ways of ensuring that implementation plans do not simply provide short term or local fixes for previous incidents. Tracking must continue until they are satisfied that revised procedures and components are seamlessly integrated into existing working practices throughout an organisation. As those procedures and components change, it may be necessary to revise previous recommendations and again track any consequent changes to ensure the continues safety of an application process.

12.4 Summary

This chapter has argued that recommendations are made in response to the causal factors that are identified by incident investigators. Some organisations, including the US Air Force [794], have argued that each recommendation must be related to a causal factor and that every causal factor must be associated with a recommendation. If a causal factor is not addressed then there is a possibility that a potential lesson will not be learned from a previous failure. If recommendations are not associated with causal factors then these is a danger that spurious requirements may be imposed for reasons that are unconnected with a particular incident. We have, however, pointed to alternative systems in which recommendations can be derived from a collection of causal factors. This often happens when investigators identify an incident as part of a wider pattern of previous failures.

This chapter has also identified a range of techniques that have been developed to help investigators derive the recommendations that are intended to prevent the recurrence of future failures or the 'realisation' of near-miss incidents. The 'perfectability' approach is arguably the simplest of these techniques. Given that many accidents and incidents are not the result of equipment failure, this approach focuses almost exclusively on the human causes of an incident. Recommendations are intended to perfect the performance of the fallible operators. An increasing number of researchers and practitioners have spoken out against this technique by arguing that investigators must focus on the context in which an error occurred [701, 342]. Instead they propose a more organisational view of failure that focuses recommendations on 'safety culture'. They have certainly provided useful correctives to the 'perfectability' approach. However, the backlash against 'prefectability' has often neglected the pragmatics of situations in which operators and managers assume some responsibility for their actions.

Subsequent sections reviewed the use of heuristics to guide the development of recommendations. These heuristics guide investigators away from interventions that are explicitly intended to rectify specific instances of human error. They also provide useful guidance on the presentation and format of potential recommendations. For instance, we have cited heuristics that encourage investigators not to propose additional studies. Such recommendations often defer actions that are then not taken when the results of additional research are not acted upon. Similarly, other heuristics are intended to ensure that investigators consider what a recommendation is intended to achieve and who must implement it.

A limitation with the heuristic approach is that it leaves considerable scope for individual differences to affect the detailed interventions that are proposed in the aftermath of an incident. Enumerations and recommendation matrices have been developed to ensure some degree of consistency between the findings of different investigators. For example, US Army publications provide lists of commonly recognise causal factors. The same documents also enumerate potential recommendations [796] These can be linked into matrices so that investigators can identify a number of potential recommendations that might be used to address a particular cause. Unfortunately, this approach only provides high-level guidance about potential interventions. The entries in a recommendation matrix tend to be extremely abstract so that they can be applied to the wide range of incidents and accidents that might be reported to complex and diverse organisations, such as the US Army. In consequence, a number of more detailed accident prevention models have been developed. These are generic only in the sense that they provide a high level framework for the drafting of proposed recommendations. The intention is that investigators can refine them to a far greater level of detail than is, typically, achieved in recommendation matrices. The barrier model has been described in previous chapters as a causal analysis technique. The same approach can also be used to guide the identification of proposed recommendations.

An important limitation with all of the approaches that have been summarised in the previous paragraphs is that they can be used to identify recommendations but not to assess their relative importance or priority. This is a significant issue for the safety managers who have to justify the allocation of finite resources in the aftermath of an incident or accident. In particular, they must ensure that the greatest attention is devoted to those hazards that are most likely to recur and which pose the greatest threat to the safety of an application. A number of proposals have been made to address these problems. Most of these attempt to synthesise incident analysis and risk assessment techniques. Many practical and theoretical problems are raised by this synthesis. we have illustrated those problems using the US Army's five stage process of risk analysis: identify hazards; assess hazards; develop controls and make risk decision; implement controls; supervise and evaluate. Previous paragraphs have described how the first three stages can be used to prioritise recommendations in terms of the difference between an initial risk assessment and the residual risk associated with both the particular causes of an incident and the more general hazards that an incident helps to identify.

The residual risk that motivates the promotion of a particular recommendation will only be achieved if the remedial actions are effectively implemented. The closing sections of this chapter show how difficult it can be to validate the claims that are implicit within a risk assessment and how hard it is to ensure conformance with recommended interventions. For example, we have brief examined the problems of documenting recommendations so that others can understand precisely what is intended and why it should be proposed. We have also looked at the difficulties of ensuring that accepted recommendations are implemented in good time across the many different operating units of complex organisations.

This closing sections of this chapter have stressed the importance of tracking recommendations. It is important to obtain feedback about how remedial actions are being implemented throughout an organisation. We have argued that implementation officers must guard against non-compliance and the deliberate violation of proposed interventions. Equally, they must ensure that the relavent personnel are provided with access to the information that is necessary to implement a recommendation. They must also ensure that this information is presented in accessible format that is easily understood by those who must use it. The following chapter examines these presentation issues in more detail. It not only considers how individual operators can be informed about the recommendations that are intended to avoid future incidents. It also addresses the more general problems of structure, format and dissemination that must be addressed when drafting incident reports. In contrast, Chapter 15 considers some of the problems that arise when investigators and safety managers must gain an overview of the many previous incidents that can motivate sustained interventions in safety-critical applications.

Chapter 13

Feedback and the Presentation of Incident Reports

This book has argued that incident reporting systems can play a prominent role in the detection, reduction and mitigation of failure in safety-critical systems. Previous chapters have reviewed a number of elicitation techniques. These are intended to encourage operators to provide information about near-miss incidents and about the failures that affect their everyday tasks. We have also identified the primary and secondary investigation techniques that must be used to recover necessary information about these incidents. This information can be used to reconstruct the events leading to failure. These models, in turn, help to drive causal analysis techniques. Finally, we have described how each cause of a failure must be considered when drafting the recommendations that are intended to avoid, or mitigate the consequences of, future failures. None of this investment in the analysis of adverse occurrenses and near-miss incidents would provide any benefits at all if the findings from an investigation cannot be communicated back to the many different groups who have a stake in the continued safety of an application process.

13.1 The Challenges of Reporting Adverse Occurrences

A number of problems complicate the publication of information about near-miss incidents and adverse occurrences. Investigators must ensure that documentation conforms both to national and international regulatory requirements. These constraints are better developed in some industries than they are in others. For example, the Appendix to ICAO Annex 13 contains detailed guidance on the format to be adopted by incident reports within the aviation industry [384]. A title must be followed by a synopsis. The synopsis is followed by the body of the report which must contain information under the following headings: factual information; analysis; conclusions and safety recommendations. Further guidance is provided about the information that should be presented under each of these sub-headings. Annex 13 also provides detailed instructions on the procedures to be adopted when disseminating the final report into an incident. This approach can be contrasted with the guidelines provided by the International Maritime Organisation's (IMO) code for the Investigation of Marine Casualties and Incidents adopted under Assembly Resolution A.849 [387]. This provides detailed guidelines on the investigatory and consultative process that must precede the publication of any report. It says almost nothing about the format and content of any subsequent documentation.

The lack of national or international guidelines provides investigators with considerable flexibility when they must document their findings about particular failures. It also creates considerable uncertainty amongst those safety managers who must ensure 'best practice' in the operation and maintenance of reporting systems. This uncertainty is well illustrated by the UK National Health Service; risk managers are responding to calls to introduce incident reporting systems without guidance on the form that those systems should take. In consequence, a vast range of local initiatives have been started to develop appropriate formats that might be used to disseminate information about adverse occurrences. The result is that hospitals have developed diverse approaches that both reflect local needs and which also make it very difficult to identify potential similarities between related incidents in different trusts. Further problems are created by the development of different national standards that can cut across these local initiatives. For example, the Royal College of Anaesthetists has taken a leading role within the UK National Health Service by issuing detailed guidance on how to gather data about critical incidents [715]. Unfortunately, the lack of national guidance in other areas of the healthcare system has resulted in standardised formats being used within certain areas of healthcare but not within others.

National and international standards are intended to support the exchange of information about previous failures. The recipients of these documents can have confidence that they will contain the information that is necessary to inform any subsequent intervention. Later paragraphs will return to the problems of encouraging this dissemination of incident reports between organisations that are often seen as 'natural' competitors. For now, however, it is important to recognise the pragmatic problems that arise when attempting to draft minimum requirements for the formatting of incident reports. The problems that arise when attempting to apply causal taxonomies and recommendation matrices illustrate how hard it can be to anticipate the nature of future failures. Similarly, it can be very difficult to predict what form an incident report should take beyond the generic and extremely abstract categories that are proposed by the ICAO.

There is a tension between the need to encourage consistency between reporting formats and the importance of allowing some flexibility in the reporting of individual incidents. The diverse nature of near-miss incidents and adverse occurrences has many further consequences of the drafting and dissemination of incident reports. As we have seen, natural language is most often used to describe the sequence of events leading up to a potential failure. The same medium is used to represent the detailed causal analysis that will, eventually, support particular recommendations. Natural language has the benefits of accessibility and flexibility. No specialist training is required to understand it. It can also be used to capture diverse aspects of an incident and its causes. Unfortunately, it can also be ambiguous and vague about key aspects of an incident. It can also be difficult to follow the large number of concurrent events that often characterise technological failure. Detailed timing issues are not well represented and it can be difficult to form coherent natural language accounts from the individual analysis of multi-disciplinary experts. The flexibility of natural language can be used to capture many different aspects of an incident. This flexibility is also a weakness because it supports the variety of interpretations that can lead to potential ambiguities. Subsequent sections will explore each of these issues in more detail.

13.1.1 Different Reports for Different Incidents

Previous chapters have emphasised the diverse nature of incidents within many industries. At one extreme, they include low-consequence near-misses that border on process improvements rather than safety issues. At the other extreme, reports provide information about high consequence failures that cannot easily be distinguished from accidents rather than incidents. This diversity has an important effect upon the nature of the documents that are used to disseminate the findings of an investigation. For example, high-consequence failures are typically reported using a highly structured format in which reconstruction is followed by analysis, analysis is followed by recommendations and so on. This formal style of presentation can be illustrated by the US Coast Guard's table of contents for a report into the loss of a fishing vessel [833]. What we have termed the reconstruction of the incident is contained within the 'finding of fact' section. The causal analysis is partly contained within these pages but is focussed on the 'Conclusions' section:

Executive Summary	3
Hearing witnesses	4
Finding of fact	6
Background of People Key to the Investigation	6
Description of the Fishery	8

Chronology	8
Description of the Vessel	28
Details of Vessel Surveys and Repairs	40
Stability Information	50
Vessel Management	66
External Environment	74
Conclusions	77
Recommendations	96

A similar format can be seen in the Australian Transportation Safety Board's (ATSB) Marine Safety Investigation reports [52]. As with the previous US Coast Guard example, the table of contents reflects the detailed investigation that was conducted in the aftermath of the incident. The reconstruction of the incident is contained within the 'Narrative' sections. Causal analysis is presented under 'Comment and analysis'. There are some differences between this report and the one described in the previous paragraph. Rather than presenting specific recommendations, the ATSB investigators identified contributing factors in the 'Conclusions' section. The lack of proposed interventions in part reflects the nature of the incident. The report's conclusions identified specific procedural problems that contributed to the incorrect loading of this particular vessel. It can, therefore, be argued that the wider publication of such specific recommendations would have had marginal benefits for a more general audience. It also reflects recent initiatives by the ATSB to move away from a 'perfective' approach towards a more 'contextual' form of analysis:

Summary	1
Sources of information	2
Narrative	5
Sun Breeze	5
The incident	5
Loading at Bunbury	6
Sailing from Bunbury, the list and subsequent events	9
Comment and analysis	15
Evidence	15
The charter party	15
Stability of the vessel	15
Stability at first departure	16
Righting lever and heeling arm curves	17
Sun Breeze: Lightship KG	19
The Class Society	20
Cargo stowage	21
Notice of intention to load timber deck cargo	22
Stowage factors, cargo weights and the masters responsibility	23
Deck cargo lashings	23
Two scenarios for the incident	25
Conclusions	27
Submissions	31
Details of Sun Breeze	33
Attachment 1	31
Stability terminology and principles	35

The task of providing feedback from incident reporting systems is complicated by the different formats that are used to disseminate information about different types of incident. For example, the previous tables describe the formal structure that is typically associated with incidents that either did, or might have, resulted in high-consequence failures. The level of detail included in the analysis is indicative of the resources that have been invested in the investigation. In contrast, many less 'critical' incidents are summarised by less formal reports. There are other reasons for exploiting a range of formats. For instance, in many industries it can be difficult to persuade operators and managers to read what are perceived to be long and complex documents about previous incidents that may, or may not, have particular relevance for their daily activities. In consequence, investigators often publish abbreviated accounts in a more 'accessible' format. They summarise the events leading to the failure and provide a brief causal analysis in two or three paragraphs. For example, the UK Marine Accident Investigation Branch (MAIB) uses its Safety Digest articles to provide a brief overview of previous incidents. Is possible to identify sentences that relate to the reconstruction of an incident, to the findings of a causal analysis and to particular recommendations. The formal distinctions that are reflected in the section heading of the more exhaustive documents, illustrated by the US and Australian reports, are not used in these summaries:

> "The ro-ro cargo/passenger ferry SATURN was completing berthing operations alongside a pier at Gourock. Prior to rigging the gangway, it was normal practice for a seaman to throw the safety net ashore from the gangway gateway, which was normally secured in the open position by hooks but, on this occasion, was not. As the net was thrown ashore, part of it became entangled in one of the gates which caused it to close and knock the seaman off balance. He was caught in the net and fell overboard, landing heavily on a pier timber before falling in the water. The seaman surfaced and, with the assistance of another crew member, managed to hold onto a pier timber. Both were recovered from the water by a fast rescue craft." [514]

This account provides a 'vignette' or 'failure scenario'. It describes an incident in an extremely compact manner. Minimal information is provided about the more detailed contributory factors that are considered in more formal reports. The previous summary does not explain the reasons why the door was not secured. In contrast, it provides readers with a direct account of the catalytic events that led to the incident and, most importantly, it illustrates the potential consequences of such incidents. Such accounts have strong similarities with the 'war stories' or anecdotes that are an important means of exchanging safety-related information within teams of operators. This is an important strength of such immediate accounts. It can be argued, however, that the lack of more sustained analysis may limit any long-term effect on system safety.

The previous paragraph provides a relatively simple example of the use of incident vignettes. Several regulatory agencies have developed variations on this approach. Investigators can use these techniques to inform readers about specific safety issues. For example, there is a danger that short vignettes will focus on the specific events that lead to a particular incident. It can then be difficult for readers to identify the more general safety issues that affect the operation or activity that was affected. There is even a danger that the readers of such incident scenarios will forget other safety issues by focusing on the specific failure described in the report. In consequence, some incident reporting systems use vignettes as a form of hook that is used to motivate readers to consider more general safety issues. This can be illustrated by a US Coast Guard report into a particular incident involving a group of sea kayakers:

> "(they) unexpectedly encountered strong currents that resulted in three kayakers being separated from the group and set out to sea. While their friends were set offshore, the main group was able to land their kayaks on a small island. Because a member of the group now ashore carried a signal mirror, the group was able to attract the attention of persons on the mainland, who in turn notified the Coast Guard. Based upon information from persons ashore, an intensive 5 hour effort was launched that eventually located and recovered the missing kayakers. This incident underscores the need for proper planning and signaling equipment, and revealed some of the inherent difficulties in mounting open water searches for objects as small as sea kayaks." [829]

The final sentence in this quotation presents the particular conclusion or finding that can be drawn from this specific incident. It also illustrates the way in which such recommendations can reveal a great deal about the intended readership of the report. The vignette is clearly not intended for the members of the rescue service. If this were the case then some additional detail should be provided about the "inherent difficulties in mounting open water searches...". In contrast, the recommendation is clearly intended for kayaking enthusiasts and to recreational sailors. After drawing this specific conclusion, the Coast Guard report goes on to reminder the reader of a number of more general safety precautions that should be followed when kayaking. The investigators place the specific recommendations about how to prevent this particular incident within the wider context of voyage planning and preparation. This is an extremely powerful technique. The particular circumstances of the incident act as a direct and clear example of the potential consequence of failing to follow safety information. It is doubtful whether the list of safety recommendations would have had the same effect if they had been presented without the incident as a preface:

> "*Voyage planning:* When planning a voyage, no matter how short or simple you intend it to be, take a few minutes to leave a float plan, including departure/arrival times, number of people and color of kayaks with a responsible friend. If it's a spur of the moment trip, write a plan just before you go and leave it in an envelope marked "FLOAT PLAN" on the dashboard of your vehicle. Make sure to always monitor the weather before and during your trip.
>
> *Know your limitations:* You alone are the best judge of your own physical limitations, the capabilities of your kayak, and most importantly, your ability to operate your craft and gear. Respect the indiscriminate power of the sea along the exposed Maine coast, and carefully avoid operating in restricted visibility, including fog, rain, and darkness..." [829]

We are currently working on a number of studies that intend to determine whether or not such presentation techniques have an impact upon decision making and risk-taking behaviour. A host of methodological problems affect such investigations. It is difficult to identify a procedure to demonstrate that individuals would be more likely to follow the safety guidance if they had been informed about previous incidents. These issues will be addressed more directly in the closing sections of this chapter when we look at the problems of validating the 'effectiveness' of incident reports.

13.1.2 Different Reports for Different Audiences

It is important to emphasise the diverse nature of those groups that have an interest in the findings of an incident investigation. Other investigators must read the reports of their colleagues to encourage consistent analysis and common recommendations to similar incidents. This also helps to sensitise individuals to emerging trends within an industry. Designers and developers may also be concerns to read incident reports in order to ensure that previous mistakes are not replicated in future systems. The operators of the application processes that are described in an incident report must also be able to access the recommendations that emerge from previous failures. This not only helps them to understand any proposed revisions to their working practices, it also helps to disseminate information about the consequences of previous failures and the potential for future incidents. These reports must also be disseminated to the managerial staff who supervise end-user activities. In particular, safety managers must be informed of any recommendations. They are often required to ensure the implementation of proposed changes. Regulators have an interest to track individual incidents. This is important if they are to monitor the safety record of individual firms. Such information helps to guide the dissemination of best practice across between companies in the same marketplace. It is also important from regulators to monitor the changing nature of incidents across an industry if they are to identify potential patterns of failure. These comments apply to national regulators. There have also been a number of international attempts to compare incident data from different countries, such as the IMO's work to collate incident reports.

National regulators and international bodies are not the only groups that are interested to learn about the insights provided by incident reports. The general public are often concerned to read the findings in these documents. This interest is often motivated by concerns over personal safety issues, including consumer protection and healthcare provision. As we shall see, many investigation agencies have responded to this concern by placing information about past failures on publically accessible web-sites. This wider interest is also being driven by an increasing willingness to engage in litigation In consequence, legal practices are often concerned to follow the incident reports in several industries. It is difficult to determine whether this public concern has been created by media interest or whether media interest has been fuelled by the engagement of this wider audience. In either case, it is important to acknowledge that all forms of the broadcast media and publishing have an active interest in reports of previous incidents and accidents.

The diverse nature of the potential readership of an incident report creates problems for those who must draft incident reports. Different reporting formats offer different levels of support for particular tasks. For example, operators are likely to require precise summaries and detailed guidance on how to meet particular recommendations. Safety managers are likely to require more information about what a recommendation is intended to achieve and how to demonstrate conformance with its particular requirement. Investigators and lawyers are concerned to understand the reasons why certain causes were identified. They may also be concerned to ensure that recommendations provide appropriate defences against any recurrence of specific causal factors. Designers and regulators will, typically, require a higher degree of technical detail than system operators. Those responsible for the implementation of future systems must also be able to generalise from specific failures to anticipate whether similar problems might affect proposed designs.

Table 13.1 illustrates the range and diversity of reports that can be generated by a single institution. This summarises the reporting activities conducted by the Hong Kong Marine Department [366]. Many of these publication requirements relate to more serious incidents and accidents. It is important, however, to see the presentation and dissemination of less critical incident reports within this wider context of regulatory and judicial requirements. The range of documents that must be produced in the aftermath of an adverse occurrence or near-miss also imposes considerable logistical demands upon such organisations. For instance, primary and secondary investigations may generate interim reports that are intended to warn operators and supervisors of any short-term actions that might help to avoid or mitigate any recurrence of an incident. These documents are, typically, superseded by the final incident report that presents the outcome of the reconstruction, causal analysis and recommendation techniques that have been described in previous chapters. As we have seen, these reports trigger implementation advisories of various forms. These guide operators and managers on the actions that must be taken to fulfill particular recommendations. Finally, statistical summaries can be derived from databases of individual incident reports. These summaries may motivate issue-based reports that investigate a number of similar incidents.

It is difficult to underestimate the logistical challenge that is posed by the production and dissemination of these different documents to the diverse groups that have an interest in a adverse occurrence or near-miss incident. For example, the recipients of an initial notification about short-term corrective actions must be informed of any longer-term measures. If this is not the case then groups and individuals may continue to exploit stop-gap measures to prevent the recurrence of previous failures. Safety managers must have access to updated statistical information if they are to determine whether or not a newly reported incident indeed forms part of a wider pattern. If this seems to be the case then they may have to obtain access to information about on-going enquiries into these related failures. Similarly, it is important that the people who contribute incident reports should receive updated information about the various levels of intervention that have been triggered by their observations.

These distribution problems must, typically, be solved within short time-limits. It is important for information to be disseminated in the aftermath of an incident. There is an obvious need to provide guidance on any short-term corrective actions. There is also a need to prevent any rumours that might be generated in the aftermath of an incident. Even in anonymous systems, it can be necessary to warn operators about the potential for future failure and to publicise the results of any secondary investigation. Conversely, it is important that any preliminary publications should not be premature. Unsubstantiated speculation can create confusion when subsequent incident reports are forced to contradict previous statements about the potential causes of an incident. These complexities not only affect the safety managers and incident investigators who must combat the causes of future incidents. They also affect the tasks of press officers and media relations officers. It is important that initial releases about an incident should not affect the results of any subsequent investigation.

Investigation Type	Summary of Process	Reporting Requirement
Informal Inquiry	Carried out by investigation board into less serious accidents. Director of Marine accepts or rejects the report and institutes follow up action.	No formal report is prepared but findings may be included in a 'Summaries' page on a web-site.
Preliminary Inquiry (Hong Kong Registered)	Director of Marine appoints professional officer(s) to conduct investigation. Report by the appointed officer is submitted to the Secretary for Economic Services with Director of Marine's observations and action. Secretary for Economic Services accepts or rejects the report.	If Marine Court is not ordered, usually the Preliminary Inquiry report will be published. If Marine Court is appointed then it reports to Chief Executive who can approve its publication and may accepts or rejects findings/recommendations of the court.
Preliminary Inquiry (Pilotage)	For minor incidents, Pilotage Authority appoints a Board of Discipline. Board of Discipline may then recommend a caution, written warning, a downgrade on pilot's licence or that Board of Investigation be held. For serious incidents, Pilotage Authority commissions a Preliminary Inquiry. This can result in Board of Investigation.	Board of Investigation submits report to Pilotage Authority. The Pilotage Authority decides on the report's recommendations and decides whether the report should be published.
Local Marine Inquiry (In Hong Kong waters)	The Director of Marine orders a Local Marine Inquiry for incidents occurring in Hong Kong waters, to be conducted by professional officer(s). The appointed officers submit a report to the Director of Marine. The Director accepts or rejects the findings/recommendations.	Findings are published as a report or as a summary on the Department web-site
Industrial Accident	The Marine Industrial Safety Section investigates incident involving repairs to any vessel, break up of a vessel, cargo handling etc.	The Investigating Officer submits report to Director of Marine for serious or fatal accidents only. Director of Marine accepts or rejects findings/recommendations. No formal report is prepared in most cases.
Conduct of Fitness Inquiry	Director of Marine can initiate inquiry into conduct of Hong Kong certified officer for "unfitness, misconduct, incompetence or negligence whether or not an accident has occurred" [366]. Inquiry is conducted by a judicial officer. Person conducting the inquiry may cancel or suspend certificate of competency/licence or censure the holder.	A report is made to the Director of Marine.

Table 13.1: Accident and Incident Reports by the Hong Kong Marine Department

It is also important not to provoke immediate calls for action without careful consideration about the justification and potential risks associated with precipitate intervention.

As we have seen, a number of different reports can be made about the same incident. Preliminary reports must be revised in the light of a secondary investigation. Final reports are informed by any subsequent causal analysis but their recommendations must be revised as regulators reassess the utility of any interventions. Figure 13.1 presents an annotated flow-chart of the procedures that support the New Zealand Transport Accident Investigation Commissions analysis of maritime incidents [631]. The rectangles that are drawn with a double line are used to denote the various publication activities that form part of a single incident investigation. These include the formal delivery of the final report. They also include the distribution of preliminary drafts to the various individuals and groups that have an interest in the outcome of any investigation.

Figure 13.1 extend the New Zealand process model with timing information. The investigation should begin within twenty-four hours of an incident being reported. The safety commission that oversees all investigations should receive a preliminary report within three days and so on. The details of such estimates depend on the nature of the incident being investigated. As we have seen, high risk incidents may justify the allocation of additional investigatory resources. In general, however, such diagrams are useful because they provide a working schedule for investigators and regulators. They also provide an important overview for operators and even for the general public who may be keen to receive feedback about the course of an investigation. In order to validate such timescales, it is important for investigators to assess the amount of time that must be spent at each stage of the analysis. Few organisations take this as far as the US Federal Railroad Administration who have estimated that is should take two hours to write an employee confidential letter, five and a half hours to review each employee statement, five hours to devise a monthly list of injuries and illness and so on [233]. The key point is, however, that if timescales are published then there must be some means of determining whether or not they are met. If they are routinely missed then either additional resources must be committed to an investigation or more realistic timescales must be published for the various participants in an investigation.

13.1.3 Confidentiality, Trust and the Media

A host of social and contextual concerns also affect the dissemination of information about incidents. Confidentiality is arguably the most important of these issues. Many organisations are concerned to ensure that reports are only disseminated to those groups that are perceived to have a 'legitimate' interest in their contents. Operators and managers are encouraged to read incident reports while strenuous efforts are made to prevent the press, lawyers and even regulators from accessing the same documents. The sensitive nature of many incidents has also created situations in which organisations are willing to sacrifice some of the potential benefits from a reporting system in order to ensure that information about previous failures is not disclosed to these 'unauthorised' sources. The ultimate examples of this sort of behaviour involve companies destroying incident databases to ensure that lawyers cannot detect examples of previous failures that might indicate negligence in failing to prevent subsequent incidents. Less extreme measures include the use of computer-based access control mechanisms that restrict those documents that a user of the system can view without specific permissions.

Some industries have also suffered from a variant of the confidentiality and security concerns, mentioned in the previous paragraph. They have become the victim of 'spoof' incident reports that are intended to undermine public confidence in their products [98, 104, 107]. These are often created by disaffected employees, by competitors or by individuals with moral and political objections to particular industries. Within an organisation it is possible to exploit a range of technologies to ensure that an incident report has been produced by an authorised individual or group. For instance, electronic watermarking embeds a code within a file. This code is difficult to alter without corrupting the contents of the report but can easily be read by authenticating software to ensure the provenance of the ocument. If the watermark code is derived from the date at which the file was last edited then this approach can also be used to detect cases in which a report had subsequently been edited or 'tampered with'. It is less easy to deal with spoof reports that originate from outside

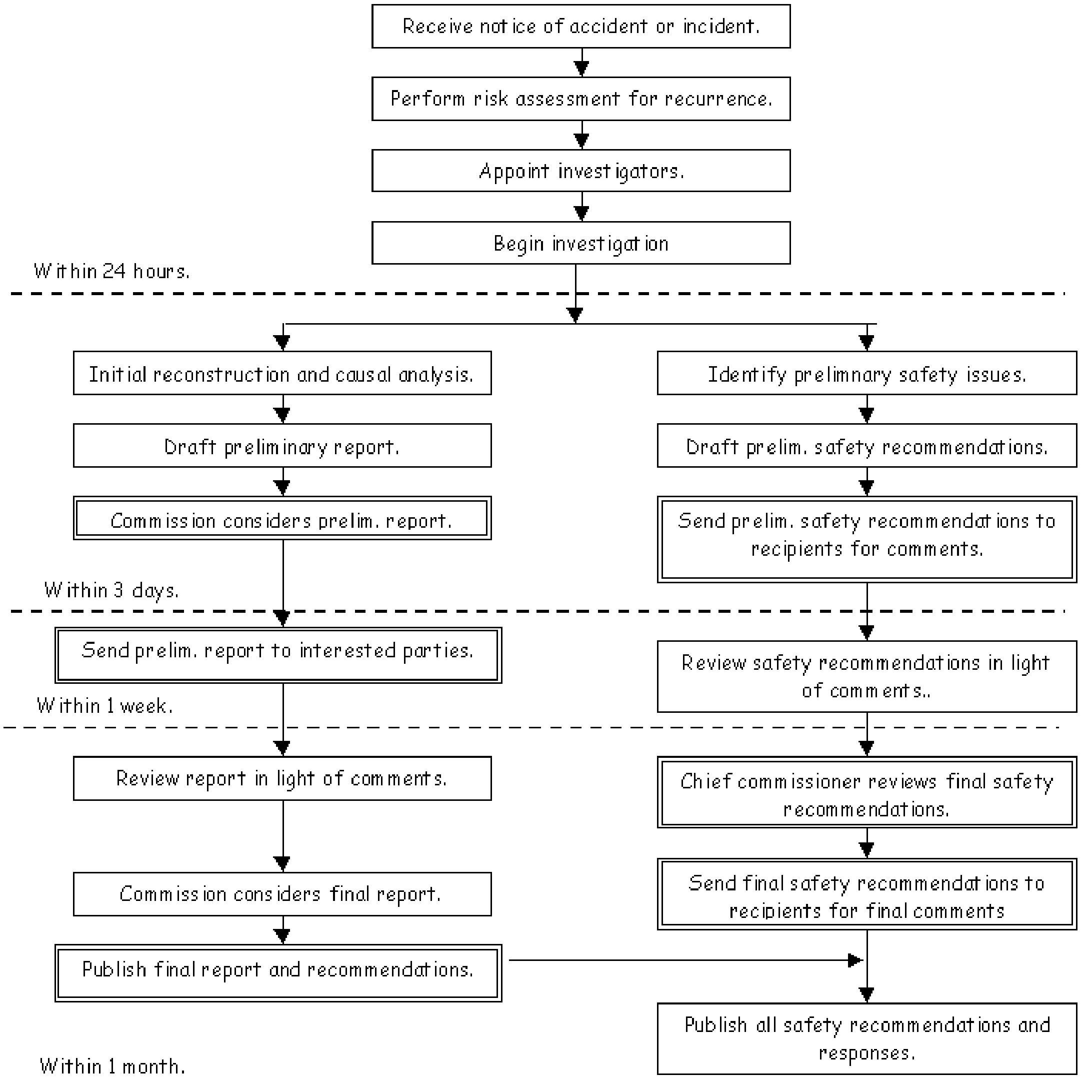
Receive notice of accident or incident.
Perform risk assessment for recurrence.
Appoint investigators.
Begin investigation
Within 24 hours.
Initial reconstruction and causal analysis.
Identify prelimnary safety issues.
Draft preliminary report.
Draft prelim. safety recommendations.
Commission considers prelim. report.
Send prelim. safety recommendations to recipients for comments.
Within 3 days.
Send prelim. report to interested parties.
Review safety recommendations in light of comments..
Within 1 week.
Review report in light of comments.
Chief commissioner reviews final safety recommendations.
Commission considers final report.
Send final safety recommendations to recipients for final comments
Publish final report and recommendations.
Publish all safety recommendations and responses.
Within 1 month.

Figure 13.1: Simplified Flowchart of Report Generation Based on [631]

an organisation. In particular, it can be difficult for safety managers and public relations staff to respond to requests for information about incident reports that they know nothing about. This creates particular problems given the concerns to preserve confidentiality, mentioned above. Denials that an incident report has been produced can be interpreted as an attempt to cover up potentially damaging information.

Many incidents involve more than one organisation. Air Traffic Management incidents often stem from the interaction between different national systems. Similarly, maritime incidents can involve ships that are registered by different Sates. Each can independently produce reports into the same incident. In consequence, national and international regulators typically require some form of coordination that is intended to encourage agreement before a final report is disseminated to its intended recipients. This is illustrated by items 4 and 5 in the following section from the Marine Accident Investigators' International Forum's Code for the Investigation of Marine Casualties and Incidents. This code was adopted by IMO Assembly Resolution A.849 (2.0). The State conducting an investigation should invite other "substantially interested" States to:

1. "question witnesses;
2. view and examine evidence and take copies of documentation;
3. produce witnesses or other evidence;
4. make submissions in respect of the evidence, comment on and have their views properly reflected in the final report;
5. and be provided with transcripts, statements and the final report relating to the investigation." [387]

Unfortunately, such high-level guidelines provide little direct help if States are forced to resole any differences over the analysis of an incident.

The difficulties of drafting and disseminating incident reports are further complicated when these documents can contain commercially sensitive information. Previous sections have argued that previous failures provide important learning opportunities. There may be clear commercial benefits to be gained from not sharing these lessons with rivals in the same market place. Similar concerns can be seen within military 'lessons learned' systems where new insights about previous failures can provide direct operational benefits. In consequence, national and international initiatives often rely upon regulatory intervention to ensure that safety-related information is disseminated as widely as possible. This laudable aim raises further questions about the format and presentation of the information that is to be shared. Participation is such schemes often implies that local systems have to conform to the minimum data standards that ensure the consistency and integrity of the common dataset. At best, these national and international presentation requirements can be integrated into existing local formats. There are, however, instances when these wider requirements are perceived to impose unnecessary additional burdens [423]. It can also be argued that these requirements reduce the effectiveness of local systems if they prevent investigators from tailoring the presentation of particular information to their immediate audience. In consequence, many systems maintain multiple versions of an incident report. An internal version can be developed to provide readers with detailed information about the particular local circumstances that contributed to a failure. There may also be a more generic account that is provided to national and international regulators. These accounts supplement the aggregated statistical data about incident frequencies that are described in Chapter 15.

Several problems can arise from attempts to maintain 'separate' accounts of the same incident. Firstly, it can be costly to support the production, distribution and maintenance of these different versions of a report. This can involve the duplication of validation activities to ensure that each account conforms to different local and national requirements. There are also additional costs associated with the archiving and retrieval of each report. As we shall see, this can involve the development of two separate but linked information management systems. Secondly, it can be difficult to ensure that these separate accounts are consistent. Even if different accounts are maintained for the best of reasons, there may still be a suspicion that internal reports are 'clean-up' or 'sanitised' before being distributed more widely. This can be illustrated by incident reporting across European Air Traffic Management systems. In one example, the manager of a national reporting system knew that a colleague in a neighbouring country had been involved in the analysis of a high-criticality air proximity violation. He was then surprised to see that they did not report any high-criticality incidents in their annual returns to EUROCONTROL. Their colleague later demonstrated that the incident did not fulfill the requirements that EUROCONTROL publish for such high-criticality mishaps. National safety managers had increased the level of criticality associated with the event because it was perceived to offer a number of key insights for the systems operating in that country. Such examples illustrate how inconsistencies between local and national or international reporting systems can arise from the best of intentions. There are other instances in which they reflect the deliberate 'manipulation' of safety-related information.

This section has briefly introduced some of the problems that complicate the presentation of information about previous failures. Chapter 15, in contrast, looks at the complexities that arise when investigators must conduct statistical analyses of aggregate incident data. In contrast, the remainder of this chapter looks at some of the existing and proposed solutions to these problems that affect individual incident reports. The analysis is structured around three generic issues that affect all reporting systems:

- *how to structure the presentation of an incident report?*
  As we have seen, there are national and international guidelines on the information that should be included within an incident report. These guidelines are not, however, available for many industries. When they are available, for example within the field of aviation incident reporting [384], they typically only provide high-level guidance about what sections should be included. They do not provide the detailed advice that is necessary when investigators begin to draft detailed accounts of previous incident. This lack of guidance has resulted in a number of poorly formatted reports in which readers have to refer to information that is distributed across dozens of pages of analysis in order to gain a coherent overview of a particular mishap;

- it how to ensure the effective dissemination of incident reports?
  Previous sections have described how the tension between a need to distribute incident reports to the many different groups and individuals who can make use of them and the need not to jeopardise confidentiality. There is also a concern to restrict 'unauthorised' media intrusion. Other systems avoid these tensions by deliberately adopting an open distribution policy. This can create problems if contributors are reluctant to submit reports that can be seen by a broad audience. Irrespective of the overall dissemination policy that is adopted, investigators face considerable logistical problems in issuing and updating information about previous incidents. Increasingly, electronic information systems are being used to reduce the costs associated with paper-based distribution. These systems are often Internet based and come with a host of implementation issues that must be considered before such applications can effectively replace more traditional techniques. They dom, however, offer considerable benefits in terms of monitoring the rate at which incident reports are accessed by their intended recipients;

- *how to validate the presentation and dissemination of incident reports?*
  Many incident reporting systems have been criticised because too much attention goes into the elicitation of data and too little goes into the effective application of that data to avoid future incidents [701]. It is, therefore, critical that some means be found of validating the particular presentation and distribution techniques are used to disseminate the lessons of previous failures. This creates a host of problems. For example, Chapter 5 has described the Hawthorne effect that can bias the results that are obtained when users know that their actions are being observed [686]. In a similar manner, direct questions about the utility of incident publications can elicit responses that may not provide accurate information about their true value.

The importance of ensuring the effective dissemination of incident reports should not be underestimated. Unless employees are provided with authoritative information about previous failures then informal networks can grow up to exchange 'war stories'. These 'war stories' provide important learning opportunities because they often encapsulate users' experiences during adverse occurrences. Unfortunately, they often over-dramatise particular incidents [749]. They can also recommend potentially unsafe interventions that contravene accepted working practices. These informal accounts are also dangerous because they exist as a form of 'distributed knowledge' that exists outside the standard safety management procedures. There is no guarantee that all staff will be told the relevant anecdotes [342]. Nor is there any certainty that appropriate actions will be taken to resolve the underlying failures that lead to the adverse incidents that are described in these accounts.

It is important to identify the intended readers of a report before investigators decide upon an appropriate structure or format for the information that is to be presented. As mentioned in the previous paragraphs, the different recipients of an incident report will have different information requirements. Tables 13.2, 13.3, 13.4 and 13.5 illustrate how information-needs grids can be drawn up to support this analysis. A separate tabular form is produced for each participant in the investigatory process. As can be seen, we have initially focussed on regulators, executive officers of board members, safety managers and operators. In local systems, some of these tables might be omitted. Some of the duties associated with safety managers might instead be allocated to system operators and hence the information needs would have to be revised appropriately. In larger, more formal systems, it would be necessary to introduce additional tables. For example, we have already described important distinctions between the information that is required by national and international regulators. Similarly, distinctions between different types of operator might result in additional tables being introduced to reflect their differing information requirements.

Regulators	Reconstruction	Causal Analysis	Recommendations
Initial Report	An initial report of the events leading to the incident and an indication of the additional information sources to the analysed	A summary of the likely causes based on the initial report together with a preliminary account of any similar incidents. Several causal hypotheses are likely at this stage.	Any immediate measures to be taken in the aftermath of the incident.
Final Report	A detailed description of what happened during the incident together with explicit justification of that account citing the evidence to support each hypothesised event in the reconstruction.	A documented causal analysis using one of the recommended techniques introduced in previous chapters of this book. This analysis may be presented in natural language but should be supported by an appendix documenting semi-formal or formal reasoning.	A detailed analysis of the proposed recommendations describing what, when, who and how they are to be implemented (see Chapter 12).
Annual Summary	Reconstruction information may be omitted in a statistical summary or annual report. It should, however, be possible for regulators to work back from the items in the summary to the more detailed final report.	If a causal taxonomy is used, see Chapter 11, then the codes or identifiers for each causal factor should be included in the statistical returns.	If a recommendation taxonomy is used, see Chapter 12, the statistical analysis should include information about the correlation of those codes to causal factors.

Table 13.2: Generic Information-Needs Table for Regulators

Each information-needs table identifies the different documents that are used to disseminate information about an incident to a particular participant group. For instance, Table 13.2 denotes that regulators should receive an initial notification in the aftermath of an incident. They should also receive a copy of the final report and an annual summary of data about all incidents that have occurred. Of course, this is not an exhaustive list. Additional interim reports may be required in some industries. Similarly, Chapter 12 has argued that closer regulatory intervention can be required to monitor and validate the implementation of particular recommendations. These caveats illustrate the generic nature of the information contained in Tables 13.2, 13.3, 13.4 and 13.5. The rows in the table must be tailored to reflect the information needs of the particular participants that are being considered.

The columns of each table identify the information that should be included within each of the documents that are issued to a particular participant. The generic information-needs tables in this chapter reflect the distinctions that have been used to structure previous chapters. Information about the reconstruction of events is followed by a causal analysis. This, in turn, supports the presentation of recommendations. Again, however, additional columns can be introduced to reflect the more detailed information requirements that are specified in some industries. For example, Tables 13.2 to 13.5 might be extended to explicitly denote whether each document should contain information about mitigating factors or about the failure of particular barriers. Other document-specific information can also be included within these tabular forms. For example, Chapter 12 has argues that it is essential to devise a timescale for the production and delivery of information in the aftermath of an incident or accident. If this is not done then there is a danger that important safety measures will be delayed. There is also a danger that potential contributors will be disillusioned by the lack of progress in addressing safety concerns. Such timetable information can be introduced into as an additional column within an information-needs table. We have not done this because such refinements can jeopardise the tractability of the tabular format. This problem might be addressed by drawing up a different information-needs table for each document that will be provided to participants in the investigatory process.

Information-needs tables are intended to help investigators identify what information must be provided to each of the participants in an investigation. Each row of the table can be used to summarise the information that they must receive. It can also be used to explain the reasons why it is necessary to provide this information to regulators, executive officers, safety managers, operators and so on. The information that is contained in each of these tables can be used in a number of ways. For example, the simplest approach is to view each row as a specification of the information needs for a single document that is to be provided to a particular group of recipients. This technique would result in a final report being produced for regulators that was quite different from the version of the final report that is presented to executive officers. The former would focus more on the generic insights derived from the incident while the latter form might provide board members with more detailed information about particular local factors. Of course, any proposed differences between these versions of a final report would have to be approved by the intended recipients. Previous sections have mentioned the suspicions that can be aroused when the internal versions of an incident report is different from that delivered to a regulatory organisation.

An alternative application of information-needs tables is to use them to derive requirements for single documents that are intended to support different participant groups. This is done by identifying similar information needs that might be addressed within a single publication. For example, there are strong similarities between the information that ‘final reports’ are intended to provide to regulators, board members and safety managers. By collating the respective requirements into a single table, it is possible to construct a checklist that can be used to determine whether any proposed report satisfies the individual requirements of each group. If a draft report does not, for example, provide safety managers with enough information about the potential need for additional data logging techniques, then it can be redrafted to support these potential recipients. Alternatively, investigators might choose to split-off this participant and draft a separate report to satisfy their particular information needs.

No matter which approach is taken, the underlying motivation for these tables is that they focus the investigators attention on the recipient’s information needs. If these are not considered early in the drafting of a report then there is a danger that individuals and groups may be denied important feedback about the course of an incident investigation. Conversely, there is a danger that some participants may be deluged by a large volume of apparently irrelevant information. Each table is

Board	Reconstruction	Causal Analysis	Recommendations
Initial Report	Executive officers within the organisation must be informed as quickly as possible about the course of events leading to an incident or accident. They must provide any necessary recourses to support a primary investigation and must coordinate any response to the media.	The causal analysis must summarise the preliminary findings but should stress any areas of uncertainty to ensure that precipitate action is avoided.	Initial recommendations should be presented in the form of a risk assessment or cost-benefit trade-off. The most plausible worst case costs and consequences of potential future failures should be summarised. The potential interventions should be identified together with any potential adverse 'side-effects' and their likelihood of preventing recurrence in the short to medium term.
Final Report	This should provide an executive summary of the events leading to an incident together with all of the information that will be provided to the regulator so that high-levels of management can respond to questions from the regulator if necessary.	The products of a causal analysis should be summarised together with references to the methods used and documentation that was produced. This is important if strategic decisions are to be justified by a detailed understanding of the mechanisms that led to previous failures.	The final report to executive officers must include a detailed list of recommendations. They must justify the allocation of resources that are required to investigate how to achieve the objectives that are specified in the recommendations section of any report.
Implementation Updates	Refer back to final report unless new evidence has been obtained.	Refer back to final report unless new evidence has been obtained.	Detailed information must be provided about attempts to validate the successful implementation of particular recommendations. This should include information from the statistical analysis of incidents and accidents that will be provided to the regulators, see Table 13.2, but will be provided to management on a more frequent basis.

Table 13.3: Generic Information-Needs Table for Executive Officers

Safety Manager	Reconstruction	Causal Analysis	Recommendations
Initial Report	The safety manager is responsible for overseeing the secondary investigation of an incident. They must be able to determine what is initially thought to have happened so that they can direct further investigation.	Initial reports from a primary investigation provide partial insights into the causes of an incident. They are, however, critical if safety managers are to allocate appropriate analytical resources. For instance, an initial causal analysis may indicate a need to provide human factors expertise or to consult with equipment suppliers.	Safety managers must safeguard systems in the aftermath of an incident. They must, therefore, consider ways of implementing those recommendations that they consider to be warranted.
Final Report	Safety managers will not only need to know the evidence that supports elements of a reconstruction, they also need to determine whether any additional logging or tracking equipment might be required to gather additional evidence about future incidents.	Safety managers must be able to determine whether or not similar causal factors have contributed to previous incidents. They may also need to assess the effectiveness of the analysis performed by their investigators, this may imply greater access to the supporting analytical documentation that is required by other parties to an investigation.	Safety manager coordinates the implementation of recommendations and so must be able to unambiguously determine the intentions behind particular proposed interventions. They must then initiate the process of determining how to achieve the recommended objectives. The safety manager will be responsible for producing the implementation reports that are passed to executive officers, see Table 13.3.

Table 13.4: Generic Information-Needs Table for Safety Managers

intended to tailor the provision of information to the particular needs of each recipient rather than allowing the provision of information to be determined by ad hoc requests.

After investigators have identified the information needs that are to be satisfied by particular reports, it is then necessary to consider the most appropriate more or form of presentation. The 'mode' of presentation refers to the medium of transmission. Most incident reports are printed, although an increasing number are being published using electronic media. Some reports continue to be delivered orally, especially in the immediate aftermath of an incident when participant must focus their attention on mitigating actions. For example, Section 67 of the Honk Kong Shipping and Port Control Ordinance requires that the owner, agent or master of a vessel to file an oral or written report within twenty-four hours of an incident occurring [365]. The format of a report refers to the content, structure and layout of information that is delivered by a particular mode. For instance, a written report may have to be submitted using an approved form. Alternatively, national regulations may simply specify the information that is to be provided without imposing any particular requirements on the particular form of presentation. For instance, section 80 and 81 od the Hong Kong Merchant Shipping (Safety) Ordinance states that the Director of Marine must be informed of any notifiable incident. This report does not have to be in a 'prescribed format', however, the form M.O. 822 "Report of Shipping Casualty" is 'recommended' [365].

Previous chapters have reviewed a range of different formats that can be used to support the

Operator	Reconstruction	Causal Analysis	Recommendations
Initial Report	Even in confidential systems, other operators may be aware that an incident has occurred. It is, therefore, often important to briefly summarise the events that led to the incident so that short-term action can be taken to prevent future recurrence.	An initial causal analysis may also be issued with the intended beneficial effect of reducing unnecessary speculation. This may, however, be premature in most cases.	The recommendations that are made in the initial aftermath of an incident must be limited by the information that is available to investigators and supervisors. It is important that no changes should be made that increase the likelihood of other forms of failure.
Implementation Report	It is important that operators receive feedback about the eventual outcomes of any investigation. This feedback is the single most critical factor in eliciting future contributions to most reporting systems. The feedback must provide a detailed account of what happened, this inevitably involves some compromise with the need to preserve confidentiality. Detailed accounts of what happened can be used to provide operators with the 'bigger picture' of events that they may not have witnessed during an incident.	Operators often form their own view of the causal factors behind an incident. These views may be modified or contradicted by the outcome of an official report. It is, therefore, important to justify particular findings. This must, typically, be done without the use of the semi-formal or formal techniques that supported an analysis given that most operators will have no experience of those techniques.	It is essential that operators understand the implications that particular recommendations have upon their working practices. They must not only be informed of what they must do and why, they must also be informed of the consequences of non-compliance and of proposed validation activities.

Table 13.5: Generic Information-Needs Table for Operators

presentation of information about particular aspects of an incident. For example, Chapter 8 has described how plans and maps can be used to supplement event based reconstructions of the events that contribute to particular failures and near-miss incidents. The same chapter also examined a range of photorealistic and model-based virtual reality techniques that have been specifically developed to support the electronic dissemination of information about particular events. In contrast, Chapters 10 and 11 have presented a number of formal and semi-formal approaches to causal analysis. These can play an important role in justifying the findings that are made in many final reports. Current documents have often been criticised because they lack any detailed justification of their causal findings [469, 426]. Why-Because graphs, ECF charts and MORT tables might all be used to format the presentation of information with particular reports. Similarly, Chapter 12 has introduced recommendation matrices, barrier summaries and risk analysis matrices that can all be used to document proposed interventions.

There is a tendency to satisfy the information requirements that are identified in Tables 13.2 to 13.5 using the products of those techniques that were used during the course of an investigation. For example, Table 13.2 argues that regulators must be provided with a detailed causal analysis as part of a final report. If investigators had themselves used ECF analysis to identify any causal factors then it would be relatively straightforward to use ECF charts within the body of their submission to national or international regulators. This would ignore the prime injunction to consider the recipient before drafting and disseminating any incident report. Within some industries, it may be entirely appropriate to exploit this semi-formal technique both to direct and document a causal analysis. In most industries, however, it would not be appropriate to expect that regulators would be familiar with this approach. In consequence, investigators must first ask whether or not the recipient of a document might be able to use any proposed form. If the answer is no, or might be no, then that form must typically be supplemented by the use of natural language descriptions. In many situations, it can be difficult for investigators to determine whether or not participants might exploit the semi-formal and formal techniques that have been explicitly developed to support incident analysis. Similarly, it can be difficult to determine whether electronic presentation techniques provide an adequate alternative to more conventional modes. The closing sections of this chapter, therefore, describe validation techniques which might demonstrate that MORT, ECF etc can satisfy recipients' information needs.

13.2 Guidelines for the Presentation of Incident Reports

The previous paragraphs have argued that both the mode and the format of incident reports must be tailored to the information needs of the intended recipients. If those recipients have limited access to computers then there are few benefits to be gained from attempts to exploit electronic presentation techniques. Conversely, if the intended recipients' usual mode of working requires on-line support then it can be particularly frustrating for them to search through, and maintain, archives of paper-based incident reports.

13.2.1 Reconstruction

Chapter 8 has considered the use of computer-based simulation techniques to support incident reconstruction. Chapters 9 has also described how investigators can exploit a range of graphical and textual notations to model the events leading to an adverse occurrence. This section looks at how investigators can communicate the products of this analysis to the wider audiences that were identified in the previous paragraphs of this chapter. In particular, we focus on the use of prose descriptions to describe the events that lead to near-miss incidents and adverse occurrences. This decision is motivated by the fact that Chapters 8 and 9 provide a detailed overview of graphical techniques. Subsequent sections in this chapter will also focus on recent advances in the use of computer-based techniques to support these prose reconstructions.

A number of constraints limit the extent to which investigators can tailor the presentation of incident information to support the particular needs of the intended recipients. In particular, they must ensure that each report satisfies any applicable national or international requirements. This can be non-trivial. For example, the AUSREP and REEFREP Australian maritime reporting systems both exploit a message format for submitting initial reports that complies with IMO Resolution A648(16) of 19 October 1989. The initial reports feed into systems that comply with the International Convention for the Safety of Life at Sea (SOLAS) Chapter V regulation 8-1, adopted by the IMO in 1996. The format of the reports derived from these systems must comply with the more recent IMO investigatory code [387].

National and international requirements often provide detailed guidance on the information that must be included in any reconstruction of an adverse occurrence or near-miss incident. For example, the reporting of maritime in incident in the UK is covered by the Merchant Shipping (Accident Reporting and Investigation) Regulations 1999. These require that the master of a vessel must send a report to the Chief Inspector of the MAIB using the quickest means at their disposal [347]. In any event, the report must arrive within twenty-four hours of the incident taking place. These initial reports must contain: the name of the ship and the vessel number or IMO identification; the name and address of owners; the name of the master, skipper or person in charge; the date and time of the accident. The report must also state where the vessel was from and where it was bound; the position at which the accident occurred; the part of ship where the incident occurred if on board; the prevailing weather conditions; the name and port of any other ship involved; the names, addresses and gender of people killed or injured. Finally, the initial report must also provide brief details of the incident, the extent of damage and whether it caused any pollution or hazard to navigation.

There can often be several different sets of applicable national and international regulations covering the dissemination of incident reports. It is possible to identify two different types of information requirement that are specified in these documents: declarative data and incident chronologies. Declarative data provides the contextual information about a system and its working environment. These details often 'set the scene' for incident chronologies. Textual and graphical time-lines focus less on the state of the system or environment prior to an incident. Instead, they focus more directly on the events leading to particular failures. The previous paragraph illustrated how the MAIB maintains relatively high-level requirements for this declarative and 'procedural' information. Other organisations have far more detailed requirements for the information that must be provided when reconstructing an incident. For example, the following list summarises the Marine Accident Investigators' International Forum [519] requirements that have been adopted by the IMO [387].

1. Particulars of voyage:
   Port at which voyage commenced and port at which it was to have ended, with dates; details of cargo and draughts (forward, aft and midships) and any list; last port and date of departure and Port bound for at time of occurrence; any incident during the voyage that may have a material bearing on the incident, or unusual occurrence, whether or not it appears to be relevant to the incident; plan view of ship's layout including cargo spaces, slop tanks, details of cargo, bunkers, fresh water and ballast and consumption.

2. Particulars of personnel involved in incident:
   full name, age, capacity on board and details of injury; description of accident; person supervising activity; first aid or other action on board; certificate of Competency/Licence: grade; date of issue; issuing country/authority and any other Certificates of Competency held; time spent on vessel concerned and experience on similar vessels and experience on other types of vessels experience in current capacity and experience in other ranks; number of hours spent on duty on that day and the previous days; number of hours sleep in the 96 hours prior to the incident; any other factors, on board or personal, that may have affected sleep whether smoker, and if so, quantity and normal alcohol habit together with information about any alcohol consumption immediately prior to incident or in the previous 24 hours; whether under prescribed medication and any ingested non-prescribed drugs and records of drug and alcohol tests.

3. Particulars of sea state, weather and tide:
   direction and force of wind; direction and state of sea and swell; atmospheric conditions and visibility; state and height of tide, in particular, the direction and strength of tidal and other currents, bearing in mind local conditions.

4. Particulars of the incident:
   type of incident together with date, time and place information; details of incident and of the events leading up to it and following it; details of the performance of relevant equipment with special regard to any malfunction; persons on bridge, in engine room and location of master and chief engineer; mode of steering (auto or manual); extracts from all relevant ship and, if applicable, shore documents including details of entries in official, bridge, scrap/rough and engine-room log books, data log printout, computer printouts, course and engine speed recorder, radar log, etc; details of communications made between vessel and radio stations, SAR centres and control centres, etc., with transcript of tape recordings where available; details of any injuries/fatalities; voyage data recorder information (if fitted) for analysis.

5. Assistance after the incident:
   if assistance was summoned, what form and by what means; if assistance was offered or given, by whom and of what nature, and whether it was effective and competent; if assistance was offered and refused, the reason for refusal.

It is important to emphasise that this is a partial list that is intended to be applicable to all types of incidents. Additional guidelines describe the more detailed information that must also be included when reports are submitted after groundings, collisions, fires etc. Such guidelines help to identify the information that should be included when documenting the events leading to an incident. They do not, however, provide detailed guidance on the format or mode of submission. As we have seen, the UK regulations simply require that masters make their initial report by the fastest means possible. The IMO's guidelines recognise that both initial and final reports can be submitted in a range of formats to be determined by the legal requirements of each State.

The IMO's requirements not only affect the initial information that must be reported in the aftermath of an incident. They are often used to guide the presentation of subsequence documents, including the final report. This can be illustrated by the Transportation Safety Board of Canada's report into the striking of a dock with an unloading boom from a bulk carrier [786]. The reconstruction of the incident begins with a section that presents background 'factual' information in a tabular format. This is illustrated in Table 13.6. This format illustrates how information requirements, such as those proposed by the IMO, can be proceduralised. The rows of the table act as a prompt to ensure that investigators provide necessary information.

ALGOBAY	
Port of Registry	Sault Ste. Marie, Ontario
Flag	Canada
Official Number	372053
Type	Self-unloading Bulk Carrier
Gross Tons	21,891
Length	222.51 m
Draught Forward: Aft:	8.05 m 8.11 m
Built	1978, Collingwood, Ontario
Propulsion	Two 10-cylinder Crossley Pielstick (10PC2-3V-400) diesel engines; 7870 kW total. Single controllable-pitch propeller and bow thruster.
Number of Crew	24
Registered Owner	Algoma Central Marine Sault Ste. Marie, Ontario

Table 13.6: Canadian TSB Tabular Preface to A Maritime Incident Report [786]

Such tabular summaries provide the contextual details that are required by national regulations, such as the Merchant Shipping (Accident Reporting and Investigation) Regulations 1999, and international guidelines, such as the IMO investigation code. They provide an overview of the technical details that are necessary in order for the reader to understand the nature of the vessel or vessels that were involved in an adverse occurrence. They also act as useful points of reference or aide-memoires that the reader can refer to as they consider a report. Rather than having to look through dense paragraphs of prose, it is possible to go back to this initial table as a reference point for information about the vessel and her crew. The example illustrated in Table 13.6 is relatively brief. Extended versions of these summary tables have been used in incident reports for other modes of transport. For example, the UK Air Accident Investigation Branch exploits a similar approach to record the registered owner of the aircraft involved in an incident, the operator, the aircraft type, the nationality of the operators, their registration, the place, date and time of the incident [15]. They use a similar tabular format to summarise information about the operators who are involved in an incident. These tables include the sex and age of the individual, the status of any operating licence, their rating, medical certification, the start of any relevant duty period, the time of their previous rest period and so on.

Such tabular forms help to structure the presentation of information that must be present if an incident reports is to conform to particular industry guidelines. They are declarative representations.

The facts that they describe, typically, hold throughout an incident. They cannot easily be used to describe more dynamic aspects of an incident. Table 13.6, therefore, illustrates a common means of prefacing more detailed reconstructions of a near-miss incident or adverse occurrence. Natural language descriptions are used in most reports to describe the way in which particular events contributed to an incident. This can be illustrated by the paragraphs that follow Table 13.6. A section entitled 'History of the Voyage' follows this summary. It describes how the vessel departed Superior, Wisconsin, at 17:10 eastern daylight time on the 7th June 1999 carrying a cargo of 26,137 tons of coal. The opening sentences of the narrative, therefore, satisfy more of the information requirements specified in the IMO's code. Unlike the contextual information that is summarised in the tabular format, illustrated in Table 13.6, this information is integrated into the description of the incident. As mentioned, it can often be difficult to locate such necessary information in large sections of prose. It is, therefore, important that check-lists be developed so that investigators can ensure that a report satisfies national and international requirements prior to publication.

After having presented the contextual information summarised above, the report goes on to describe the events immediately preceding the incident. As with most incident reports, a discussion of more latent causes is postponed until the analysis sections. The narrative is often presented in as simple a form as possible. There is an attempt to minimise any additional commentary on the significance of particular events. This too is postponed until the subsequent sections of analysis. This is an important strength of many incident reports because the readers' interpretation of particular events is not biased by a premature commentary. Equally, however, it can be difficult for readers to determine which events in a narrative will turn out to have a critical significance for the causal analysis and which are introduced to provide additional background. For instance, the following quotation presents the subsequent sections of the 'History of the voyage' section in the Transportation Safety Board of Canada's report. It is difficult to determine whether the unscheduled maintenance will or will not play a significant role in the course of the incident until readers complete the remaining paragraphs:

> "On the morning following departure, as the vessel crossed Lake Superior, the chief engineer, with the authorisation of the master and shore management, shut down the port main engine for unscheduled repairs. At 17:45, the vessel advised Vessel Traffic Services that the passage through the locks at Sault Ste. Marie and the St. Marys River would have to be conducted with one engine. Permission was granted by the United States Coast Guard and Sault Ste. Marie harbour master. Before arriving at the Sault locks, the vessel had developed a one and one-half degree list to port. Concerned with low water levels in the St. Marys River and the 10 cm of extra draught the list would have created, the master ordered the second officer to lift the unloading boom and slew it to starboard as the vessel departed Poe Lock at 0225 on June 9. In doing this, it was hoped that ballast water remaining on board in the No. 3 port ballast tank would be displaced to starboard and the list corrected. As the boom was being lifted from its saddle, it began swinging to port. Despite attempts by the second officer to check its movement with the slewing controls, the boom accelerated outwards until it contacted the front of the accommodation at an angle of 90 degrees to the vessel." [786]

At first sight, it would appear that the unscheduled maintenance is no more than a contextual detail. Investigators mention it in the report because it represents an unusual event that took place before the problems with the boom. In subsequent paragraphs, however, the reader learns that as soon as the boom began to hit objects on the shoreline the master put the engine to 'full astern'. Unfortunately, this overloaded the single remaining engine. The watchkeeping engineer, therefore, informed the master that he would have to reduce power. The master eventually arrested the vessel's forward motion by ordering the lowering of the stern and both bow anchors. This illustrates the manner in which the unscheduled maintenance removed a potential barrier, the main engine, that might have prevented further damage once the incident had begun. This argument is not explicitly presented in the incident report until the analysis section. In consequence, it is very easy for readers to overlook the significance of the maintenance event even as they read about the problems of putting the remaining engine 'full astern'. As mentioned, this technique avoids prejudicing the reader. By separating the reconstruction and the analysis, individuals are encouraged to form their own hypotheses before reading the investigators' interpretation. On the other hand, this approach can be deeply frustrating. During recent interviews, a safety managers referred to the 'Perry Mason' or 'Agatha Christie' style of incident reporting. The reader never understands the true significance of a particular event until they read the final pages of a report. As a result, they must re-read the report several times in order to identify the way in which the elements of a reconstruction contribute to a causal hypothesis.

A Structural Analysis of Incident Reconstructions

There are a number of different ways in which investigators can structure their presentation of the events leading to an incident. For example, they can use prose accounts to summarise the elements of graphical time-lines such as those introduced in Chapter 9. This approach describes events in the order in which they occurred during an incident. Alternatively, investigators can exploit a more thematic approach in which the time-line leading to a failure is described from one particular perspective. Subsequent paragraphs then go back to the beginning of an incident to present the flow of events from another perspective. These different approaches have a number of strengths and weaknesses. For example, readers can easily trace the ordering of events if they are described in the order in which they occurred. This can, however, create a false impression. The reader is presented with a global view of all the events that occurred across a diverse system for each point of time. The participants in an incident, typically, would not be in such a fortunate position. Further problems affect the use of more thematic approaches. If investigators describe the course of events from particular perspectives then it can be difficult for readers to piece together an overview of the concurrent failures that often characterise many incidents. The following list, therefore, summarises these different techniques for structuring the presentation of a prose reconstructions:

- *a single chronology*;
  As mentioned, this approach simply reconstructs the time-line leading to an incident. Each significant event is described in chronological order. This provides a relatively simple overview of an incident. An important benefit of this approach is that readers can quickly scan the text to identify what events occurred at any particular point in time. This scanning is facilitated by using time-stamps as marginal notes that can act as indices into particular paragraphs:

  08:00 repairs were begun; however, while removing the cylinder head from the engine, it was discovered that a cylinder head stud was broken and would have to be replaced.

  08.30 (est.) The chief engineer, who had previous experience with this type of repair, informed the master of the situation and revised his estimated completion time upwards to 24 hours.

  ... ...

  20:00 the head tunnelman reported to the chief engineer that he had finished raining and cleaning the hydraulic system. The chief engineer indicated that, be cause the head tunnelman was unfamiliar with the procedure to bleed the air from the hydraulic system, he decided to wait until daylight the following morning.

  The limitation with this approach is that events from many diverse areas of an application process can be listed next to each timestamp. This can create problems because these entries will not be uniformly distributed over time. This implies that for any particular system, there may be relevant information scattered across many dozens of paragraphs. This imposes considerable burdens upon readers who want to piece together what happened to a particular subsystem or operator.

- *a single chronology with backtracking*;
  A number of further problems affect the use of single chronologies to structure the presentation of any incident. As we have seen, catalytic and latent events are not uniformally distributed across the time-line of an incident. Typically, initial failures may lie dormant until a number

of triggering conditions defeat any remaining barriers. This creates problems because it can be difficult for readers to gain an overview of an incident. An initial description of the latent failures can quickly become swamped by the mass of details that typically accompany any presentation of catalytic failures. Many investigators have responded to this problem by starting the reconstruction of an incident with a brief overview or summary of the events leading to a failure or near miss. Subsequent paragraphs then go back to the start of the catalytic failures to examine those events in greater detail. The Transportation Safety Board of Canada's report exploits this approach. The initial summary, presented in previous paragraphs, is followed by a more detailed reconstruction of the engineering and maintenance activities during the incident:

> "At 0800 on June 8 the repairs were begun; however, while removing the cylinder head from the engine, it was discovered that a cylinder head stud was broken and would have to be replaced. The chief engineer, who had previous experience with this type of repair, informed the master of the situation and revised his estimated completion time upwards to 24 hours. As the work would be conducted near the running starboard main engine, the chief engineer suggested that the vessel be stopped in Lake Superior for the duration of the repair. In consultation with the master and chief engineer, the company engineering superintendent decided that the vessel should proceed towards Sault Ste. Marie in case further shore support was needed for the repairs. During the previous sailing season, the vessel had operated for several months on one engine, including during passages through the American locks at Sault Ste. Marie..." [786]

This is intended to provide readers with the contextual framework, including latent failures and mitigating factors, that is necessary to understand the significance of particular catalytic events. Unfortunately, this approach also suffers from a number of limitations. In particular, the use of a more detailed chronology for catalytic failures can reduce the amount of attention that is paid to latent failures. This potential bias is often countered by reports that devote most of the subsequent analysis sections to the longer-term causes of an adverse occurrence or near-miss incident.

- *multiple thematic chronologies*;
  Previous chapters have attempted to distinguish between incident reconstruction, which explains what happened, and causal analysis, which explains why an incident occurred in the manner that it did. These distinctions have been maintained because they are, typically, reflected in the structure of most incident reports. An initial discussion of the events leading to a failure are then followed by a discussion of the causes of those events. These distinctions can, however, become blurred in some reports. For example, some accounts are structured around several different chronologies. These time-lines each reflect a particular analytical approach to the incident. For example, an account of the human factors failure may precede a description of the events that contribute to any system failures. The subsequent analytical sections in the incident report are then used to 'weave' together the individual events that are included in these different chronologies. An explanation of systems failure may be given in terms of human factors issues, or vice versa. This approach has much to recommend it. The individual chronologies can be used to demonstrate that analysts have considered a suitable range of potential causal factors before performing their analysis. The causal analysis, in turn, provides an explicit means of unifying these disparate accounts. There are, however, a number of problems with this approach. It often makes little sense to provide a chronology of operator actions without also considering the system behaviours that they were responding to and were helping to direct. Their are logistical problems in ensuring consistency between these multiple chronologies. It can also be expensive to recruit and retain the necessary technical expertise to construct these different perspectives, especially in small-scale local systems.

- *multiple location-based chronologies*;
  A variation on the previous approach is to present different chronologies that record the events taking place within particular locations or subsystems during an incident. For instance, the

report can describe the events on the bridge before describing what happened in the engine room. This approach provides readers with some impression of what happened to individuals and systems within that particular location. It avoids the false 'global' view that can often makes readers wonder why operators did not intervene to rectify what to them is an 'obvious' problem. There is, however, no guarantee that readers will avoid this potential pitfall even if location-based chronologies are used to structure a report. Each successive account contributes to their understanding of an incident. The cumulative insights that can be obtained from reading each of these accounts would clearly not have been available to operators or line managers. There are also more pragmatic problems. It can be difficult to ensure that these different accounts are consistent with each other, especially when materials and other forms of communication pass between different locations. A common flaw in this form of incident report is to find that a message has been sent from one location but that its receipt is never mentioned in subsequent descriptions. In such circumstances, the reader cannot easily determine whether the message was never received or that it did arrive and the investigator simply omitted to mention its receipt in their reconstruction.

This is a partial list, investigators have used a number of hybrid techniques that draw from several elements of this list. Many reports also combine textual chronologies with some of the graphical and diagrammatic reconstruction techniques that were introduced in Chapters 8 and 9. Subsequent sections of this chapter will describe how these combined approaches have recently been combined to support the on-line publication of incident reports. For now, however, the key point is that investigators must consider the consequences that prose chronologies have upon the intended recipients of the report. If an extended single chronology is used then investigators can help readers to navigate a reconstruction by providing timestamps as marginal indices and by using different paragraphs to describe concurrent events in different areas of the system. If investigators do not consider the potential weaknesses of these formats then there is a danger that the resulting document will fail to support the various user groups that have been identified in previous paragraphs.

The task of reconstructing an incident does not simply depend upon the chronology that is developed. Investigators must also determine what to include and what to omit from any reconstruction. The following list summarises potential guidelines that might be used for determining what information should be included in a reconstruction:

1. *is the information required by regulators?*
   Bodies such as the IMO enumerate the information that must be provided in many incident reports. These requirements typically focus on declarative data about the type of system that was involved in an incident. It can be difficult to identify a suitable format with which to present this information. The statistical nature of much of this data lends itself to tabular formats rather than the prose descriptions that are used in other sections. They also focus on gathering information about the catalytic events leading to a failure.

2. *is the information necessary to understand catalytic events?*
   Chapter 10 has shown how ECF analysis can proceed by reconstructing the flow of events back from the point at which energy was 'transferred'. Conversely, investigators might use P-Theory to work forward from the first event that deviated from the normal pattern of operation. In either case, there is a focus on the catalytic events that contributed to an incident. It may seem to be relatively straightforward to present this material. The previous paragraph has, however, summarised the problems that arise when concurrent interactions can contribute to the course of an incident.

3. *is the information necessary to understand latent events?*
   We have also argued that it is important to understand the longer-term factors that contribute to an incident. For example, Chapter 3 described how systems may not be in a 'normative' state for many years. For example, working practices can evolve to remove important barriers. This creates considerable problems for the investigators who must determine how best to present this material. If a linear chronology is used then it may not be easy for readers to understand how an apparently insignificant event contributed to an eventual incident. The

significant of that description may only emerge many dozens of pages later. Alternatively, if backtracking is used then the latent events can be described together with more immediate 'triggering' conditions. As we have seen, however, such techniques can provide a perspective that was not available to operators at the time of any failure.

4. *is the information necessary to support the narrative of other events?*
   Some events are included not because they are essential to the readers' understanding of an incident but because they link other more important events. These events can create considerable confusion. For instance, many investigators maintain the distinction between analysis and reconstruction by separating them into different chapters of a report. In consequence, it can be difficult for readers to distinguish between these 'filler' observations and latent or catalytic events. They can, therefore, help to spark alternative causal hypotheses that must be explicitly rejected in any subsequent analysis if readers are to be satisfied by the investigators' interpretation of events.

5. *is the information necessary to eliminate certain hypotheses?*
   One means of restricting the number of putative hypotheses that might be evoked by any reconstruction is to explicitly provide information about events or conditions that were not apparent during an incident. For example, investigators often begin a reconstruction by providing information about the prevailing weather conditions. If they were bad then readers are informed of a potential cause of any failure. This information is, however, also included for incidents that occur under favourable conditions so that readers can better interpret the subsequent chronology. Such techniques must be used with care. It can be argued that in seeking to inform the readers' interpretation of key events, investigators may be introducing an unwarranted bias into their accounts of an incident.

6. *does the information describe a failed barrier?*
   Investigators often decide to provide information about the protection mechanisms that were intended to protect a system and its operators. These events often stand-out from reconstructions because they describe how cross-checks were not made. They may also describe decisions that were contradicted or countermanded to achieve particular operational goals. As with the previous items in this list, the presentation of such information can create certain problems for those who must draft incident reports. Many chronologies describe these checks without explicitly stating that this was an opportunity to protect the system from a potential failure. This is an appropriate approach because the reconstruction of what happened is separated from the causal analysis of why it happened in that way. Equally, however, it can lead to disjointed accounts where catalytic events are interrupted by accounts of apparently insignificant conversations between key personnel. It may then take many pages before readers learn that these conversations might have prevented or mitigated the consequences of the incident.

This is a partial list. Chapters 8 and 9 present further requirements for the information that must be considered by any reconstructions. These requirements can also inform the presentation and dissemination of information to the readers of incident reports. For instance, investigators must extend the scope of any reconstruction to include remedial or mitigating actions. We have not extended the list to explicitly include these items because they have already been addressed in the previous chapters. In contrast, the items in this list describe the issues that must be considered when presenting particular elements of a reconstruction. For instance, it can be difficult for readers to understand the role that a latent failure can play in an incident if it included in a reconstruction without any supporting explanation. The elements in this list also help to highlight a number of more general issues. For example, investigators may decide to separate analysis from reconstruction in the manner recommended by the previous chapters of this book. This does not, however, imply that readers will simply switch off their analytical skills as they read a reconstruction and then switch them back on again as they start the section labelled 'analysis' or 'findings'. For example, it is tempting to interpret any situation in which an individual questions the safety of an operation as a failed barrier. A reader who forms this belief while reading the reconstruction of an incident may retain this impression even if subsequent sections do not consider the event any further and if the operation has little influence on the outcome of an incident.

A Case Study in the Literary Criticism of Incident Reports

The following paragraphs use a report that was drafted by the ATSB to illustrate some of the issues raised in the previous paragraphs. This case study was chosen because the investigators exploit a simple single chronology with limited backtracking. It, therefore, exploits a relatively simple narrative structure. s will be seen, the form of analysis that is applied to this incident report resembles the techniques that are used in literary criticism. This is entirely intentional. It is important to recognise the prose techniques that investigators often implicitly exploit when drafting their reconstructions of adverse occurrences. As we shall see, these techniques often have an important impact upon the readers of an accident report. The report concerns a stability problem that affected a merchant vessel, the Sun Breeze. The reconstruction begins with a textual summary of the declarative information that was provided Table 13.6 in previous incident reports:

> "The Panama flag Sun Breeze is a 11,478 tonne deadweight general cargo vessel, owned by NT Shipping SA. It was built in 1998 by Miura Shipbuilding Co Ltd of Japan and is classed with ClassNK (Nippon Kaiji Kyokai). The vessel was delivered to the owner on 9 February 1999, six months prior to the incident. Sun Breeze is 109.30 m in length overall, has a beam of 19.8 m and a summer draught of 9.264 m. Propulsive power is delivered by a seven-cylinder Akasaka diesel engine developing 5,390 kW driving a single fixed-pitch propeller and providing a service speed of 13.5 knots..." [52]

This declarative summary continues by describing the dimensions of the vessel's holds and its ballast capacity. It also described the previous expertise of the crew. For example, the master 'had sailed as mate for five years on general cargo ships, log carriers and bulk carriers and had 22 years experience in command of various vessels, mainly bulk carriers' [52]. The reconstruction then goes on to describe how 'no untoward incidents' were reported on previous voyages between Japan and Indonesia, Singapore, Malaysia and Thailand. The report then focuses on the chronology of events leading to the incident. The master received a fax from the charterers of his vessel on the 2nd August. His voyage instructions were to load a minimum of 10,000 m^3, 'up to the vessels full capacity' of jarrah and karri... The instructions also described how there were approximately 650 packs of 4.2 metre lengths of timber, 650 packs of 3 metre lengths, 20 packs of 4.5 metre lengths and 820 packs of 6 metre lengths. The master calculated that this would require some $14,354m^3$ without any allowance for broken stowage [52].

Previous paragraphs have considered how the Sun Breeze case study satisfies the declarative requirements imposed by the IMO. The opening pages of the narrative describe the vessel, its crew, the cargo and the nature of the proposed voyage. The report then goes on to trace latent communication problems between the master and the company that was chartering his vessel. He tried to find out if they wanted to load the minimum agreed figure of 10,000 m^3? If so then the entire cargo could be loaded underdeck. If not then did they want to load the 14,354 m^3 of cargo in the instructions? The shippers acting on this presumption had prepared to load about 15,000 m^3 of cargo. This would exceed the vessel's bale capacity. The charterers replied by noting that the cargo had been fixed with them on the basis of lump sum freight and that the vessel had been accepted. This illustrates the way in which the ATSB investigators use the preliminary paragraphs to reconstruct the situation that was faced by the master of the Sun Breeze before he commenced the loading of his vessel. This is important because it enables the reader to follow the way in which the context for an incident developed from its initial stages. Readers are informed of the relationship between the charterer and the master. He voiced his concerns but was ultimately reminded of the contract that he was expected to honour. His concerns were also partly addressed by the development of a loading plan that met the charterer's requirements. It is also important to note that, although the opening sections of the chronology provide important information about the context in which the incident occurred, they only hint at the events that eventually threatened the safety of the vessel. It would surprise the reader to find that the incident did not involve the way in which the cargo was stowed on-board the Sun Breeze. Equally, however, the opening paragraphs provide few cues about the eventual nature of the incident. This narrative technique avoids the problems associated with reconstructions that provide information that could not have been available to the participants in an incident.

The report follows a simple, single path chronology. The background events, mentioned above, are followed by an account of the more immediate events that led to the incident. The Sun Breeze arrived at Bunbury at 23:24 on 15 August 1999. Subsequent paragraphs describe how loading commenced at 09:00 on the 16th August. It is important to emphasise, however, that many incident reports exploit a number of different chronologies. For example, the ATSB's reconstruction goes on to develop a multiple, location-based chronology. The report surveys the operations to load the cargo. While this was going on, the report also describes how the third mate took various steps to improve on these operations. Previous paragraphs have mentioned the burdens that this can impose upon readers. It can be difficult to reconstruct the time-line of events that led to an incident if individuals must piece together the partial orderings of multiple chronologies. This problem is often apparent in the marginal notes that are made by the readers. For example, printed incident reports often contain informal sketches and time-lines that individuals have made to help them keep track of the parallel events that are described in the sequential accounts of prose reconstructions.

Previous paragraphs have argued that readers cannot simply disengage their critical and analytical faculties as they read the reconstruction section of an incident report. In consequence, there is a continual temptation to filter and analyse the information that is presented in any account. This can be illustrated by the following paragraph from the ATSB's report. At one level, it describes how the ship's master managed to convince the harbour master that he knew how to address any stability concerns. At another level, the same prose can be interpreted as describing the failure of a barrier that might have prevented the incident. Previous paragraphs in this chapter have described a number of reasons why investigators might include particular information in an incident report. It is important to recognise that a careful reading of such documents will yield not only the basic event structure in any reconstruction but also the investigators intentions behind the presentation of particular information:

> "Between 06:00 and 09:00 that morning, after loading had been completed, the harbour master noted that the ship initially had a list to starboard of about 4 degrees. It then had port list of about 4 degrees before becoming upright. He became concerned about the vessels stability... The harbour master went on board Sun Breeze at about 15:00 to discuss the vessels stability with the master. The ship was upright then and the harbour master recalled the master saying that he was going to transfer fuel oil from high tanks to low tanks to provide additional stability. The master gave the harbour master a copy of the stability calculation for the vessels departure condition, indicating that the GM, after correction for free surfaces effects in tanks, was 47cm. The harbour master asked the master how he knew what the weight of the cargo was and whether the packs of timber were marked with weights. The master said that the packs were not marked with weights and that he had estimated cargo weights by draught survey. The masters reply gave the harbour master the impression that the master knew what he was doing." [52]

Investigators have clear intentions when they introduce such narratives into incident reconstructions. They describe how a potential barrier, provided by the harbour master's stability checks, were circumvented. The question that this paragraph raises is whether or not the intended audience for a particular report would be able to identify this intention as they read the narrative reconstruction. The previous sections of this book have introduced accident and incident models that provide semantics for terms such as 'barrier', 'target', 'event', 'condition' and 'catalytic failure'. These concepts form part of a vocabulary that many readers of an incident report will not have acquired. This has important consequences. Some readers will be able to interpret the intention behind particular elements in a reconstruction. Other readers of the same document may view them as random observations that seem to contribute little to the overall report. These individuals may require considerable help in the subsequent analysis if they are to filter the mass of contextual information and 'filler' events to identify key aspects of an incident. The closing sections of this chapter will describe a range of validation techniques that can be used to determine whether or not such differences can jeopardise the effective communication of safety information within an incident report. As we shall see, however, these techniques have not been widely applied and there is little empirical evidence about individual differences in the interpretation and analysis of incident reports.

One of the key problems in reading an incident reconstruction is that it can be difficult to distinguish key events from the mass of background detail that investigators often include in their accounts. As mentioned, these background details include contextual information that is necessary to establish the circumstance in which an incident occurred. They also include the less important 'filler' events that link together other more significant aspects of an incident. This can be illustrates by the ATSB's description of the events that occurred immediately after the Sun Breeze left Bunbury. The first sentence can be described as a filler; 'When the harbour master disembarked at 18:15, he returned ashore and drove back to the wharf where Sun Breeze had been berthed, watching the ship'. The second sentence provides important information that must be supported by interviews that were conducted after the incident had occurred; 'he had some lingering concerns about the vessels stability but, when there seemed to be no problems as the vessel proceeded outbound, he returned home'. Readers face a number of further problems in anticipating what will and what will not turn out to be key elements of any reconstruction. For example, the same paragraph in the report continues; 'the vessel was being set to the east by the tide and he adjusted the course to 335 degrees, using about 5 or 10 degrees of port rudder to do so'. It is difficult to determine how important the bearings will be for the subsequent analysis of the incident. An initial reading cannot determine the significance to attach to the change of course. One consequence of this is that the reader of an incident report often have to read reconstruction sections several times after having read a subsequent analysis in order to understand how particular events contributed to an eventual incident [749].

The previous paragraphs in this section have presented a *structural analysis* of the ATSB report. For instance, we have shown how this document exploits several different chronological structures. A single linear thread can branch into consecutive accounts that each depict parallel events in different locations. It has been argued that this can impose significant burdens on the reader of an incident report who must piece together these accounts in order to derive an overview of the events leading to a near miss or adverse occurrence. Similarly, we have identified some of the problems that can arise from the usual practice of presenting a reconstruction before any causal analysis. Some readers may be forced to re-read narrative accounts several times before they can place key events within the time-line of an incident. We have also argued that it can be difficult to entirely separate analysis from reconstruction. A careful reading of an incident reconstruction will not only provide information about the 'flow' of events, it can also reveal the investigators' intentions behind the presentation of particular events. This structural analysis should not obscure the importance of the prose that is used in any reconstruction. Investigators must tailor their use of language so that readers can clearly follow the flow of events. It is important that the prose should not over-dramatise the incident by adding literary effects that do not contribute to the exposition. The ATSB avoid this potential pitfall and provide a valuable example of the precise and concise use of language to describe what must have been an 'extreme' situation:

> "The 3rd mate changed back to manual steering and ordered 10 degrees of port helm to bring the vessel back on course. At this time, the vessel started listing to port. The mate, who was on the bridge at the time, told the 3rd mate to telephone the master. The masters phone was busy so the 2nd mate went below to call him. When the mate was on his way to the bridge to take over the watch, he noticed that the vessel was taking a port list. He thought that the vessel might have taken a 15 degrees list but, by the time he got to the bridge, the vessel was coming upright again. He telephoned the master asking him to come to the bridge, after which the vessel took a starboard list. When the master returned to the bridge the list was about 15 degrees or 20 degrees to starboard. He stopped the engine. The rudder was amidships but the vessel was still turning to starboard. The list continued to increase as the vessel turned slowly. The ship attained a maximum list of about 30 degrees to starboard before it reduced to about 25 degrees.

> At about this time, lashings on the cargo on no. 1 hatch top released when securing clips opened and nine packs of timber were lost over the side." [52]

The prose that describes subsequent events also exhibits this sparse but effective style. It also provides further examples of the way in which a single chronology will branch at key moments during an incident. In particular, the report uses consecutive parallel chronologies to describe the remedial actions on shore and on-board the vessel. The basic structure is further complicated by the use of consecutive time segments. In other words, an initial paragraph describes how a distress broadcast was received by a volunteer group who, in turn, triggered the response on-shore. A subsequent paragraph then describes the crews' actions while the police and harbour master were coordinating their efforts. A further paragraph then resumes the account of the shore-side activities. This technique is similar to the way in which movie directors frequently cut between parallel streams of 'action'. It is, however, a difficult technique to sustain throughout the many pages of prose narrative that can are presented in many incident reports. As in many films, the different strand of activity that we have termed 'chronologies' are brought together by joint efforts to resolve the incident: The harbour master drove to the pilot boat and eventually helped to coordinate the use of the tug Capel to disembark some of the Sun Breeze's crew. A single chronology is then resumed as the report describes how the master reduced the list to about 5 degrees by altering the ballast in the vessels tanks. A surveyor joined the Sun Breeze and performed more detailed stability calculations. These identified problems both in the previous assumptions that had been used by the crew and in the factors that had been included in their calculation. The vessel eventually berthed again at Bunbury, where the cargo was secured and some of the packs were removed; 'the vessel sailed at 2005 on 25 August for the discharge port in China, arriving there at 0600 on 10 September 1999 without further incident' [52].

Less Detailed Reconstructions...

The directness of the prose style that is used in the ATSB report is even more important for incident reports that summarise less critical failures. These documents may be limited to a few brief paragraphs that must not only reconstruct the events that led to the incident but must also document any analysis and summarise the subsequent recommendations. Space limitations are not the only constraint that complicates the task of drafting these less formal accounts. The national and international requirements, listed in previous pages, do not apply to the less 'critical' failures that may only be reported to internal company schemes. Greater diversity is permitted for incidents which are perceived to offer a relatively low risk from any potential recurrence. A report into a minor workplace injury need not record the position of the captain and the chief engineer, as required under the IMO guidelines that were cited in previous paragraphs. Similarly, the information that is required in any reconstruction is also tailored according to the audience. The exhaustive lists of information requirements compiled for the UK MAIB and the IMO are intended to ensure that national and international regulators can access necessary details. This amount of contextual information is often irrelevant to the masters, operators and employees who must endeavour to avoid future incidents. Investigatory agencies, therefore, do not include all of the informations that is provided in a final report to them when they disseminate accounts of an incident or accident to the industry that they protect. This point can be illustrated by the summary reports that various agencies have published to disseminate information about previous incidents These documents strip out much of the contextual information and what we have described as 'filler' details to focus in on particular hazards. As can be seen, four sentences are used to reconstruct the near-miss incidents. The investigators also summarise the recommendation to ensure that fuel containers are separated from potential ignition sources and secured to prevent shifting:

> "Portable space heaters are frequently utilised on-board vessels in this fishery to warm divers and to keep sea urchins from freezing. Coast Guard personnel have observed, at sea, a number of fishing vessels using these portable heaters while also carrying portable fuel containers, including those used to carry gasoline for outboard engines. If gasoline or other flammable liquids are carried on-board a vessel it is critically important to ensure

> that these items are well separated from any potential ignition sources and secured or lashed in place to prevent shifting. Accidental spillage of any flammable liquid, especially gasoline, in the vicinity of an open flame source can result in a catastrophic fire that will quickly engulf a vessel." [827]

Such brevity and directness is achieved at the cost of much of the detail that is provided in the previous example of the Sun Breeze report. Relatively little is said about the range of heaters and storage devices that were observed on the vessel. Nor do the Coast Guard state whether or not they saw any particularly hazardous uses of space heaters and fuel storage containers. It can be argued, however, that such details are irrelevant to the intended readers of this account. The Coast Guard have identified a generic problem based on their observation of previous incidents. The recommendation is also suitably generic so that individuals can apply the advice to guide their daily operations.

It might appear from the preceding discussion that the level of detail that is provided in any reconstruction should simply reflect the nature of the incidents that are being reported on. The Sun Breeze was an inherently more complex incident, involving potentially greater risks from any recurrence than the 'simpler' problems observed by the US Coast Guard. The ATSB could have summarised the previous incident in a few lines; 'various communications failure contributed to a failure to correctly load the cargo, this ultimately compromised the stability of the vessel'. Such a summary would strip out necessary information so that readers would have little opportunity to gain the many insights that the investigators derived from their more detailed reconstruction of this incident. Conversely, 20-30 pages of reconstruction could have been devoted to the reconstruction of previous incidents involving the storage of flammable substances on-board oyster boats. It is far from certain that such a detail analysis would contribute much beyond the existing summary [827].

It is, however, too superficial to argue that the nature of an incident determines the depth of any reconstruction. A number of additional factors must be considered when investigators decide how much detail should be introduced into a reconstruction. In particular, Chapter 2 identified the tension that exists between the need to provide sufficient contextual information for operators to understand the ways in which an incident occurred and the potential problems that can arise if a report threatens the anonymity or confidentiality of a submission. If too many details are provided about the context in which an incident occurred then the readers of a report may be able to infer the identify of the person who originally instigated a subsequent investigation. Conversely, if too few details are provided then operators may feel that any recommendations are not properly grounded in the detailed operational circumstances of their working practices [423]. The elements of the following list present some of the issues that help to determine the amount of detail that must be introduced into any reconstruction:

1. *what are the boundaries of trust and confidentiality?*
   There may be significant constraints upon the amount of detail that an investigator can include within any reconstruction. Participation rates in any incident reporting system can be threatened if previous assurances of anonymity are compromised by an investigators publication of particular items of information. This can lead to a difficult ethical decision in which investigators might choose not to release important information about a potential hazard in order to safeguard the longer term future of the reporting system [444].

2. *how serious is the incident?*
   If investigators are not constrained by bounds of confidentiality then the level of detail in a report can be determined by an assessment of the potential seriousness of an incident. Chapters 11 and 12 have described how risk assessment techniques can be used to assess the potential threat the might be posed by the recurrence of an incident. This can be estimated in terms of the probability and consequence associated with each of the hazards that contributed to the incident. Investigators must also account for the risk associated with those hazards that were identified during an investigation but which did not actually occur during a near miss. In general, these techniques are only likely to provide subjective estimates that cannot easily be validated until some time after a report has been published and disseminated.

3. *how complex is the incident?*
   Chapter 1 reviewed Perrow's argument that the increasing coupling of complex application processes is producing new generations of technological hazards. These hazards are being generated faster than techniques are being developed to reduce or mitigate those hazards [675]. If this is correct then it will be increasingly difficult to summarise an adverse occurrence or near miss incident in only a few sentences or paragraphs. The complexity and coupling of application processes defy attempts to summarise their failure. It is certainly true that incidents involving the use of high-technology systems, in general, require greater explanation than those that involve less advanced systems. It is difficult, however, to be certain that this is a result of the inherent complexity of high-technology systems or the lack of familiarity that many potential readers have with their design and operation.

4. *how much of the context can be assumed?*
   Investigators can often omit contextual information if they assume that such information is already widely known amongst their intended readership. This provides a partial explanation for the hypothesised differences between reports into the failure of high and low technology systems. Investigators need not provide additional information about the nature and use of fuel containers because they can assume that most readers will be familiar with these components. On the other hand, greater details must be provided for similar reports into the failure of navigational radar systems. Such assumptions can, however, be ill-founded. As we have seen, it can be difficult to anticipate the many diverse groups and individuals who have an interest in any particular incident report. In consequence, investigators can make unwarranted assumptions about their knowledge of particular working practices.

5. *how significant are the recommended changes?*
   Investigators may be forced to provide a detailed reconstruction of an incident or incidents in order to demonstrate that particular recommendations are justified by previous failures. If they propose radical changes to particular working practices or considerable investment in additional plant then safety managers must be confident that those recommendations are warranted by a detailed analysis of a near miss or adverse occurrence. In particular, it can be important to identify those plants, systems or production processes that have been affected by previous failures. If this is not done then operators can claim exemptions from particular recommendations on the grounds that any subsequent analysis is not based on evidence from their production processes.

6. *can details be introduced to achieve particular effects?*
   The introductory sections of this chapter argued that investigators must consider the information requirements of the various readers of an incident report. Any proposed document must support the different needs of regulators, safety managers, operators and so on. These requirements clearly have an effect on the level of detail that is required in any reconstruction. As we have seen, safety managers will require a considerable attention to detail if they are to accept radical reforms. Conversely, system operators may require less detailed information in order to motivate them to implement specific changes to their daily routines [864]. Not only can investigators vary the level of detail in incident reconstructions to support the different end users of a report, they can also adjust these details to achieve specific effects upon those readers. For example, the US Coast Guard place their warnings about the dangers of fuel cannisters in the context of the oyster fleet even though the same warning might be applied to everyone at sea. It can be argued that a generic or abstract warning 'never put fuel cannisters next to an ignition source' would have lacked the immediacy and directness of the report that places the analysis and recommendation within the context of a brief reconstruction of previous near misses involving particular vessels that *were observed* by particular Coast Guard officers.

The previous list identifies a number of factors that can influence the level of detail that investigators provide in the reconstruction of a near miss or adverse occurrence. Several of these concerns can be illustrated by a case study from the UK MAIB's Safety Digest. This provides a relatively brief account of a particular set of failures . The location of the incident and the types of vessel involved are all clearly identified in the reconstruction. Anonymity is less important than disseminating this contextual information which, as we have seen, can contribute to the authenticity and 'directness' of a report. As we shall see, however, this detailed example is then used to make some very generic points about the nature of system failure:

> "The Veesea Eagle, a 622gt standby vessel, was on station in the North Sea. Early one morning, the superintendent received a call saying that No 1 generator had failed due to an exhaust pipe failure. Shortly afterwards, he received a further report saying that because of a damaged piston, No 2 generator had also failed. Back on board, the harbour generator was started to enable repairs to No 2 generator to be undertaken. While the company arranged for a replacement standby vessel, the chief engineer started to replace the damaged piston. Meanwhile, the harbour generator also failed. Although the main engine was still functioning, and steering was available by using the independently driven Azimuth Thruster, the company decided to tow the vessel back to port for repairs." [517]

This example illustrates how relatively short reports can be well situated within particular locations and operating environments. The vessel is clearly identified as the Veesea Eagle and its purpose and methods of operation can be assumed from its role as a standby vessel in the North Sea oil fields. The previous list, however, makes the point that such assumptions can be dangerous. Some potential readers may lack the prior knowledge that is necessary to infer the style of operation from the first sentence in the report.

This MAIB report is instructive for a number of further reasons. At one level, the previous quotation describes a 'freak' collection of failures that successively denies operators of their reserve power sources. The author of this report, however, takes considerable care to use this particular incident to make several more generic points about the nature of incidents and accidents. They achieve this using a relatively simple technique that avoids much of the jargon that often weakens more academic work on accident models. These points are illustrated by the findings that follow the reconstruction cited above:

1. No 1 generator
   Had been running successfully following a complete overhaul, including a new crankshaft, earlier in the year. After the vessel re-entered service, the chief engineer adjusted the fuel timing to improve performance. When he left, the relieving chief engineer also adjusted the fuel timing, but had not been told about the last adjustment. The result of the latter was massive after burning damage to a piston head, cylinder head, and exhaust trunking.

2. No 2 generator
   Had also been running successfully, when a piston failed for no apparent reason.

3. Harbour generator
   Failed because of a lack of lubricating oil and it wasn't monitored. [517]

There is no mention of a Swiss cheese model or of dominoes or of latent and catalytic causes. The analysis does, however, provide clear examples of the ways in which different types of failure combine to breech redundant defences. It is tempting to argue that the format and layout of the report reflect the investigator's intention to explain how each defence was breeched. This implicit approach might, in time, encourage readers to recognise the value of this style of analysis. Such an analysis is not as unrealistic as it might sound, given that readers are often expected to infer the more general lessons that can be derived from such specific reconstructions. Unfortunately, there is no way of knowing whether this was a deliberate intention on the part of the investigator or whether they simply described the reasons why each power source failed on demand.

13.2.2 Analysis

The previous paragraph illustrated the way in which the MAIB presented the particular findings that were derived from the Veesea Eagle. The following paragraphs build on this analysis to identify a number of more general issues that must inform the presentation of any causal analysis in incident reports. This chapter draws upon case study reports that have been published by maritime investigation authorities in countries ranging from Australia to Japan, from Hong Kong to Sweden. Before reviewing the presentation of causal findings in these reports, it is important to emphasise that the case studies reflect existing reporting practices within the maritime domain. A number of important characteristics differentiate these practices from those that hold in other areas of incident reporting, for example within commercial aviation or the medical industries. Maritime incident reporting systems have some advantages over these other systems. In particular, as we have seen, the IMO provides a structure for ensuring some degree of consistency between the incident reporting systems that are operated in member states. Equally, there are a number of features of the maritime industry that create problems which are less apparent in the aviation domain. For example, many maritime occupations are characterised by a plethora of relatively small companies. The structure of the in-shore fishing industries provides a strong contrast with that of the global market in commercial passenger aviation. In consequence, although we focus on particular case studies that are drawn from the maritime industries the following analysis will also occasionally digress to look at wider issues that affect the presentation of causal analyses in wider domains.

Very few maritime incident investigations exploit any of the semi-formal or formal causal analysis techniques that have been introduced in Chapters 10 and 11. This is not to say that causal techniques have not been proposed or developed for these industries. For example, Pate-Cornell has demonstrated how a range of existing techniques might have been applied to the Piper Alpha accident [665]. Wagenaar and Groeneweg's paper entitled 'Accidents At Sea : Multiple Causes And Impossible Consequences' exploits a variant of the Fault tree notation introduced in Chapter 9 [852]. Unfortunately, this pioneering work has had little or no impact on the reports that are disseminated by investigation agencies. The situation is best summarised by recent proposals for a research program to 'tailor' causal analysis techniques so that they can match the requirements of the maritime industries. The US Subcommittee on Coordinated Research and Development Strategies for Human Performance to Improve Marine Operations and Safety under the auspices of the National research Council helped to draft a report advancing a 'prevention through people program' [835]. This report proposed that in order for the Coast Guard to interpret information about previous successes and failures 'a methodology using root cause (or factor) analysis tailored for the marine industry and taking legal issues into account, could be developed'. Such 'a system of analysis' would 'trace the chain of events and identify root causes or factors of accidents in maritime transportation system'. The general lack of established causal analysis techniques within the maritime industry is also illustrated by very similar European research initiatives. For instance, the European Commission recently funded the Casualty Analysis Methodology for Maritime Operations (CASMET) project [154]. This was to facilitate 'the development of a common methodology for the investigation of maritime accidents and the reporting of hazardous incidents, to improve the understanding of human elements as related to accidents and account for these aspects in the common methodology'. Although the project developed a high-level architecture for such a method it did not have the impact on regulatory organisations that was anticipated. As we shall see, the CASMET criticisms of previous practices could equally be applied today as they were when the project was launched in 1998. They argued that 'present schemes are rooted in a compliance culture, in which the competence and focus are by tradition oriented towards guilt-finding' [154]. Even those countries that have established independent investigation units will still initiate legal proceedings if an investigation reveals a violation. They further argue that this inhibits individuals from contributing to an incident reporting system. This quotation reflects the results of a survey that the CASMET project conducted into existing incident and accident investigation techniques within the European maritime industries. The emphasis was on identifying situations in which individuals and organisations failed to comply with regulations. Few attempts were made to perform any form of deeper causal analysis to understand why such failures occurred. Their analysis raises important issues but it is slightly superficial. Chapter 5 has analysed the limitations of 'no blame' systems in comparison to the 'proportionate blame' approaches that have been adopted within areas of the aviation and healthcare communities. There are a number of further reasons why maritime incident reports are not guided by the causal analysis techniques that have been applied in other domains. One important consideration is that statutory requirements often focus on the findings that are derived from a causal analysis rather than the process that is used to produce them. The use of particular techniques is, typically, less important than that any incident should be investigated and reported on within particular time limits. For instance, the UK Merchant Shipping (Accident Reporting and Investigation) Regulations 1999 simply state that the master will 'provide the Chief Inspector with a report giving the findings of such examination and stating any measures taken or proposed to prevent a recurrence' within fourteen days of a serious injury [347].

Such requirements do not simply reflect a concern with identifying violations, they also reflect the need to address important safety concerns within a specified timelimit. They also reflect a pragmatic desire to maximise limited investigatory resources. These are particularly stretched by the transient and dynamic nature of the maritime industry. It is an obvious point but the location in which an incident occurs does not remain in the same position as it might do in many other industries. Indeed, it may move outside of the territorial waters of the nation in which the incident occurred. This gives rise to the National Research Council's concern that any causal analysis technique must recognise the particular legal concerns that characterise the maritime industry [835]. The complexity of addressing these legal issues often results in a high-level of ambiguity in the international recommendations that are intended to guide both the conduct and the publication of any causal analysis. For instance, the IMO's Code for the Investigation of Marine Casualties and Incidents simply defines a cause to mean 'actions, omissions, events, existing or pre-existing conditions or a combination thereof, which led to the casualty or incident' [387]. While this indicates a relatively broad view of causation, it provides little concrete guidance to investigators who must first identify appropriate causal analysis techniques and then determine how best to publish the findings of such techniques. In contrast, the IMO code continues by identifying the high-level, organisational processes that might support any analysis. It avoids imposing any constraints on the techniques that might be used; 'marine casualty or incident safety investigation means a process held either in public or in camera conducted for the purpose of casualty prevention which includes the gathering and analysis of information, the drawing of conclusions, including the identification of the circumstances and the determination of causes and contributing factors and, when appropriate, the making of safety recommendations' [387].

The lack of specific guidance on appropriate causal analysis techniques and or presentation formats for the findings of such approaches does not imply that international regulatory organisations are insensitive to the problems of analysing incident and accident reports. The 20th session of the IMO assembly adopted resolution A.850(20) on the 'human element vision, principles and goals' for the Organisation. This stressed that the ships' crews, shore based management, regulatory bodies, recognised organisations, shipyards and even legislators contribute to the causes of near miss incidents and adverse occurrences. It was argued that "effective remedial action following maritime casualties requires a sound understanding of human element involvement in accident causation... gained by a thorough investigation and systematic analysis of casualties for contributory factors and the causal chain of events''. It was argued that the "dissemination of information through effective communication is essential to sound management and operational decisions". Unfortunately, the resolution did not identify appropriate means for reconstructing such causal chains nor did it describe effective means of dissemination. In contrast, the resolution identified a number of high-level goals that mirror the more general objectives of the IMO's investigation code. Item (e) describes the high-level objective of disseminating the lessons learned from maritime incident investigations within the wider context of safety management within these industries:

- "(a) to have in place a structured approach for the proper consideration of human element issues for use in the development of regulations and guidelines by all committees and sub-committees;
- (b) to conduct a comprehensive review of selected existing IMO instruments from the human element perspective;
- (c) to promote and communicate, through human element principles, a maritime safety culture and heightened marine environment awareness;
- (d) to provide a framework to encourage the development of non-regulatory solutions and their assessment based upon human element principles;

- (e) to have in place a system to discover and to disseminate to maritime interests studies, research and other relevant information on the human element, including findings from marine and non-marine incident investigations; and
- (f) to provide material to educate seafarers so as to increase their knowledge and awareness of the impact of human element issues on safe ship operations, to help them do the right thing." [388]

The lack of specific guidance provided by national and international organisations is understandable. Their main concern is often to ensure that incident and accident reporting systems are established in the first place. Unless incident and accident reporting systems are perceived to meet particular local needs then there is a danger that they will be under-resourced or abandoned. The imposition of detailed requirements for causal techniques or presentation formats might jeopardise the ability of individual agencies to tailor their system to meet those local needs [423]. In consequence, different nations have adopted a radically different approaches to incident analysis even though they are signatories to the same IMO resolutions. These differences have a significant impact on the manner in which the findings of a causal analysis are both presented and disseminated. For instance, the Japanese maritime incident reporting system is based upon a judicial model.

> "The inquiry takes an adversarial form pitting the parties concerned, against each other. 'Open court', 'Oral pleadings','Inquiry by evidence' and 'Free impression' are employed in the inquiry. Moreover the independence of the Judge's authority in exercising inquiry is laid down by the Marine Accidents Inquiry Law. The examinee, counselor and commissioner may file an appeal with the High Marine Accidents Inquiry Agency within seven days after the pronouncement when he is dissatisfied with a judgement pronounced at a Local Marine Accidents Inquiry Agency. A collegiate body of five judges conducts the inquiry in the second instance through the same procedure of the first instance. When a judgement delivered by the High Marine Accidents Inquiry Agency is not satisfactory, it is possible to file litigation with the Tokyo High Court within 30 days after the delivery to revoke the judgement."[393]

The Swedish reporting system is far less adversarial. The purpose of maritime investigations is to provide 'a complete picture of the event'. The intention is to understand why an accident or incident occurred so that effective preventive measures can be taken to avoid future failures. The Swedish investigation board argue that 'it is to be underlined that it is not the purpose of the investigation to establish or apportion blame or liability' [768]. Such contrasting approaches have important implications for the nature of the findings that are likely to emerge from any investigation. They can also have a profound impact upon the techniques that might be used to disseminate those findings. For instance, in the Japanese system the individuals who contribute to an incident will have direct disciplinary action taken by the investigatory board if they are found to be negligent. In contrast, the Swedish system arguably adopts a more contextual approach that seeks to understand the circumstances that contribute to a failure rather than simply punish any particular violation or error. Such distinctions are an over-simplification. They ignore many of the cultural influences that have a profound effect on the manner in which these systems are actually operated. They also ignore the detailed motivations that justify particular approaches. These caveats are important because other researchers and analysts have often been too quick to identify a 'blame culture' or a 'perfective approach' in systems that they are not involved in. Although it is clear that the Japanese system exploits a judicial process, this does not mean that it ignores the contextual factors that lead to errors and violations. Although the Japanese system may ultimately prosecute individuals and groups, this is the very reason why a judicial process is exploited:

> "Marine accidents occur frequently not merely from human act or error but from a complexity of factors like working conditions, ship and engine structure, natural circumstances such as harbor and sea-lane, and weather and sea condition. On the other hand, material and circumstantial evidence is often scant. So it is sometime difficult to grasp what actually happened and to investigate its cause. Since disciplinary punishment against the holder of a seaman and pilot's competency certificate may restrict

> their activities and right, a lawsuit-like procedure is used in marine accidents inquiry to ensure careful consideration and fairness." [393]

One cannot assess the way in which an incident reporing system is operated simply by the aims and mission statements that they espouse. Abstract comments about the multi-facted nature of incidents and accidents need not result in practices that reflect a broad view of causation. It is, therefore, important to look beyond such statements to examine the findings that are produced by these different systems. For instance, the Japanese Marine Accident Inquiry Agency have published a number of the incident reports that they have submitted to the IMO. These include an analysis of a collision between a dry bulk carrier, Kenryu-Maru, and a cargo vessel, Hokkai-Maru [391]. The reconstruction exploits two location-based chronologies. The events leading to the collision are described first from the perspective of the Kenryu-Maru's crew and then from that of the Hokkai-Maru. The Kenryu-Maru reconstruction describes how the third mate found the echo of Hokkai-Maru on the radar but did not report it to the master . The Summary of Events section continues that 'the master went down to his room without a concrete order that the third mate should report him when the visibility become worth'. Some time later, the third mate again established radar contact with the Hokkai-Maru 'but he did not set eyes on HOKKAI-MARU carefully'. He changed course 'degrees but he did not watch out for HOKKAI-MARU carefully on her radar to determine if a close-quarters situation is developing and/or risk of collision exits' nor did he use 'sound signals in restricted visibility'. The report describes how the third mate oredered the vessel hard to starboard when he eventually recognised that the echo of the Hokkai-Maru was crossing their course. He made an attempt to telephone the master 'but he could not dial as it was so dark'. The collision occurred shortly afterwards.

The style of reconstruction used by the Japanese Marine Accident Inquiry Agency in the Hokkai-Maru report is very different from that described in previous incident reports. There is a strong element of interpretation and analysis in the Summary of Events section. For instance, the reconstruction includes the statement that 'the third mate did not know (about) this dangerous situation, as he did not watch carefully by her radar... he believed each ships could pass on the port each other'. Such sentences are, typically, to be found in the analysis sections of the reports produced by organisations such as the US National Transportation Safety Board (NTSB) or the UK MAIB. In contrast, the findings from the analysis of the Hokkai-Maru incident were presented to the IMO as follows:

1. Principle findings:
   Both KENRYU-MARU and HOKKAI-MARU did not sound signals in the restricted visibility, proceed at a safe speed adapted to the prevailing circumstances, reduce her speed to the minimum at which she can be kept on her course and take all her way off if necessary, because they did not keep watch on each other by radars. Both masters did not ordered to report to them in case of restricted visibility. Also both duty officers did not report to their masters on restricted visibility.

2. Action taken:
   The masters of KENRYU-MARU and HOKKAI-MARU were inflicted reprimand as a disciplinary punishment. The duty officers of the both ships were inflicted suspensions of their certificates for one month as a disciplinary punishment.

3. Findings affecting international regulations: No reported" [391]

It can, therefore, be argued that the judicial nature of the Japanese maritime reporting system has a profound impact on the presentation of a reconstruction, on the interpretation of an adverse occurrence and on the findings that are drawn from an analysis. Similar collisions have led other organisations to look in detail at the precise communications that occurred between members of the crew immediately before the incident [826]. These findings have prompted further research into what has become known as Bridge Resource Management. Conversely, other incident reports have looked at the particular demands that are imposed on crewmembers who use radar for two potentially conflicting tasks. Such systems are typically used both for target correlation, for example to identify the position of the Hokkai-Maru, and for navigation support [406]. It should be stressed that it is impossible to tell whether these broader issues were considered during the investigation and analysis of the incident. It is clear, however, that the findings that are published in the report focus on the responsibility of the individuals who were involved in the incident.

Chapter 12 has argued that a perfective approach that focuses exclusively on individual responsibility and blame is unlikely to address the underlying causes of many incidents. This does not imply, however, that the Japanese approach will be counter-productive. We have already cited the objectives of the investigation agency. These clearly recognise the complex, multi-faceted nature of many incidents and accidents. It is, therefore, possible to distinguish the broader aims of the reporting system from the findings of judicial enquiries. It is entirely conceivable that other actions may be instigated within the transport ministry to address the broader problems that are illustrateed by the Hokkai-Maru incident even though the investigation agency publishes a narrower view that reflects a focus on individual responsibility. Such a hybrid approach avoids situations in which operators view a reporting scheme as a means of avoiding disciplinary action for their actions. It is, however, important not to base such wide-ranging observations upon the analysis of a single report. The Hokkai-Maru case study might not be typical of the other reports that are produced by this system. It can be argued that the investigator who drafted the findings in this report was unusual in their focus on individual blame. Alternatively, the crews' actions in this incident might genuinely be interpreted as negligent in some respect. A careful reading of the other submissions to the IMO does, however, confirm that this style of analysis is part of a wider pattern. The judicial nature of the investigatory and analytical processes results in findings that focus on individual responsibility for the causes of adverse occurrences and near miss incidents. For instance, another report describes the grounding that ultimately led to the loss of the Bik Don, a cargo vessel. The findings of this report again focus on the role of the master rather than the circumstances that helped to influence his actions; 'the grounding was caused from that the master did not check the circumstances of course, which he selected newly for short-cut, on charts or sailing directions' [392]. However, no disciplinary action was taken even though the vessel was lost and the crew had to be rescued by a patrol boat of the Japanese Maritime Safety Agency. There are number of puzzling aspects to this incident report. It is unclear why the analysis identifies human error as the primary cause and yet the judicial system does not take any retributive action. This might reflect a recognition that the Master had suffered enough from the loss of a vessel under their command. Alternatively, this might be an indication that the investigatory authorities had recognised some of the contextual issues that influence human error. There is some evidence for this in the report's 'Summary of Events' which again contains detailed causal analysis:

> "He thought she could pass through south of the island safely because he sometimes saw the same type ships as her navigated that area. He had no experience to pass through south of Shirashima islands. But he did not confirm about Meshima island shallow waters extended from Meshima island, and especially independent shallow water named Nakase that was apart 1,000 meters from the Meshima island shallow waters by checking charts or sailing directions before his deciding." [392]

This case study illustrates two key points about the presentation of any causal analysis. Firstly, readers can experience considerable practical problems if the interpretation of an incident is distributed throughout a report. It can be necessary to read and re-read many different chapters in order to piece together the reasons why particular causes were identified. Unless this can be done then it is difficult to justify or explain the recommendations that are proposed in the aftermath of an incident. In the previous example, the decision not to act against the Master could only be explained by piecing together elements of the analysis that were presented in the Summary of Events and in the Principle Findings sections of the IMO report. Secondly, it is important that investigators explain the reasons why particular conclusions were reached and why other potential causes were excluded during an analysis. In this case, the incident report focuses exclusively on the actions of the Master. It does not consider factors such as external pressures to meet contract obligations that motivated the navigational error. In contrast, the reader is simply informed of the finding without any of the intermediate analysis or documentation that might accompany the application of the techniques described in Chapter arefpart:analysis and 11, such as MORT or ECF analysis. It is important to note that the IMO reports, cited above, provide only brief summaries of each adverse occurrences. The documents are only between five and six pages in length. The tractability of these documents might be significantly compromised if these additional details were included. This, in turn, can act as a powerful disincentive for many readers who are daunted by the task of carefully reading incident reports that often run to dozens of pages in length. There are, however, a number of techniques that might be used to address this problem. For instance, summary reports might explicitly refer to the more detailed documents that are required to justify a causal analysis. In the Japanese maritime system this might include summary transcripts of the judicial process that is used to examine causal hypotheses. ECF charts or other more formal documentation might be provided. Alternatively, the findings of a causal analysis might be justified in a natural language summary. Clearly the choice of approach must depend on the investigatory process, the resources that are available to support any analysis and on the nature of each incident. Without this information, however, readers can be left with considerable doubts about the findings of many incident investigations [749].

The Japanese approach to the presentation of causal findings within IMO summary reports can be contrasted with the approach adopted by the Swedish Board of Accident Investigation. The high-level guidelines that describe their approach to causal analysis are very similar to those presented by the Japanese Marine Accident Inquiry Agency. They both stress the way in which incidents stem from multiple causes. As we have seen, the Japanese argue that "...marine accidents occur frequently not merely from human act or error but from a complexity of factors" [393]. The Swedish approach is summarised as follows:

> "An accident can normally not be attributed to one cause. Behind the event are often enough a series of causes. An investigation is therefore aimed at the consideration of different possible causes to the accident. Many of these will in the course of the investigation be eliminated as improbable. This elimination however means that the remaining possible causes gain strength and will develop into probable causes. Needless to say, both direct and indirect causes must be considered." [768]

There are also some interesting differences between the Japanese Marine Accident Inquiry Agency's guidelines and those published by the Swedish Board of Accident Investigation. As might be expected, the Japanese stress the need to protect those involved in an incident when the findings of any report 'may restrict their activities and rights'. A judicial process is required because it can be 'difficult to grasp what actually happened and to investigate its cause' [393]. The value of these defences cannot be underestimated in proportionate blame systems. The judicial proceedings provide for the representation of the individuals concerned in an incident. This contrasts strongly with many Western systems in which the causal findings of an investigation can be reported, often without any detailed supporting analysis, and with only a minimal participation of those involved in an incident. The Swedish system focuses less on issues of individual representation and protection. In contrast, it focuses on the more detailed problems that must be addressed by any causal analysis. They make the important distinction between direct and indirect causes. This resembles the distinction that we have made between latent and catalytic failures. Their high-level guidelines also stress the need to eliminate certain hypotheses in order to identity potential causes. Previous sections have argued that it is important that readers can understand why particular hypotheses have been eliminated if they are to have confidence in the findings of any causal analysis [199, 198].

The following paragraphs analyse a case study report that follows the Swedish Board of Accident Investigation's guidelines. This can be contrasted with the case studies from the Japanese reports to the IMO, although it should be remembered that these are summary descriptions whereas the Swedish report provides a more detailed analysis. This incident resulted in a collision between the Swedish vessel, MT Tärnsjö and the Russian vessel, MV Amur-2524 [767]. The Swedish report follows the formatting conventions that have been described in previous section. A summary section is presented which includes an initial set of recommendations. This is then followed by factual information. The 'Analysis' chapter is then presented before the conclusions and a re-statement of the recommendations. Our focus here is on the presentation of the analysis section. This begins with an analysis of the place and the time of the accident. There were "concurring statements from the pilots and the crews, the accident took place approximately 100m west or west north west of Strömskärs northern point" [767]. There was, however, some disagreement about whether the collision occurred at 18:28 or 18:30. This reinforces the point that reports do not simply focus on the causal analysis of an incident. They often also include an assessment or interpretation of the evidence upon which such an analysis is grounded. The Tärnsjö report continues with an analysis of the speed at which each vessel approached the point of collision. This again illustrates the way in which the investigators analyse the various accounts of the incident to make inferences that will eventually support their causal findings:

> "The accounts of the course of events show that the Tärnsjö passed Toppvik about the same time that Amur-2524 passed Nybyholm. The distance between Toppvik and the collision site is approximately 2.8 M. Hence the Tärnsjö sailed about twice as far as Amur-2524. The Tärnsjö's average speed on the stretch becomes 10.5 or 9.4 knots... Toward the end of the stretch, her speed was reduced by the reversing and was approximately 2 knots at the collision. This means that her speed during the first part of the stretch must have been above the average. Taking the lower average speed of 7.6 knots and assuming the full astern manoeuvre was somewhat hampered by the ice and therefore began 0.3 M before the collision site, then the Tärnsjös speed before reversing becomes 10 knots. If one assumes instead that the reversing was fully effective as in the emergency stopping trial, then the same calculation gives a speed just before the reversing of close to 10 knots. Against this background the Board considers that Tärnsjös speed when entering the yaw off Tedarö light must have been considerably greater than the 6-7 knots stated by the pilot." [767]

As mentioned, the analysis of the evidence is used to support particular findings about the course of the incident. These, in turn, support any subsequent causal analysis. In this case, the Swedish board found nothing to show that the reversing procedure was less effective than what the stopping trial had revealed. From this they concluded that the reversing procedure was probably initiated too late and that this, combined with the high speed of the Tärnsjö, prevent the crew from stopping the vessel before the collision.

Many investigators might have concluded their analysis with the finding that one of the crews had failed to apply the full astern manoeuvre in time to avoid the collision. In contrast, the Swedish report follows the distinction between direct and indirect causes made in the guidelines, cited above. The investigators go on to analyse the communications that took place between the vessel prior to the incident. The analysis section goes on to explain that the pilots of the two vessels had met on at least two occasions prior to the incident during which they discussed the passing manoeuvre. The report observed that "from then on they did not communicate nor give any information regarding positions or speeds; neither did they agree upon how they would handle the meeting, in other words, at what speed they would meet or if one vessel should heave to" [767]. The analysis then focuses on the role played by the officers of the watch in the interval immediately before the incident. As before, the investigators use the evidence that was identified during the reconstruction of the incident to support particular findings about the crews' behaviour. In this instance, it is argued that the Tärnsjös chief officer and second officer "did not participate in the navigation nor did they follow what was happening other than temporarily during the passage through the Hjulsta bends" [767]. In consequence, they were both unsure about the vessel's speed before the collision. Neither could recall when and where the reversing procedure was started. The report describes a similar situation on the Amur-2524. The helmsmen and the master did not participate in the navigation once responsibility for this had been handed to the pilot. After having analysed the evidence that was obtained for the events immediately before the incident, the analysis section then goes on to consider more proximal factors. It prefaces this by summarising the outcome of the investigation which draws upon the analysis of the evidence that was presented in previous paragraphs:

> "The immediate cause of the collision was that the meeting was poorly planned. This in turn was mainly due to faulty communication between the pilots regarding how and where the meeting would take place. The Board however is obliged to note a circumstance that it has had reason to mention in several previous investigations; that is, the

> tendency that vessels with a pilot on board are conned by the pilot alone without use of the other resources available on the bridge. The accident could have had far worse consequences had the Tärnsjö's stem collided with the Amur-2524 a little further aft and also penetrated the athwartships bulkhead to hold nr 3. If this had happened Amur-2524 would probably have capsized and sunk." [767].

This quotation not only illustrates the way in which the Swedish report gradually build up a causal analysis from the investigators' interpretation of the available evidence. It also makes explicit the investigators' assessment of the worst plausible consequences of an incident. Chapters 5, 7and 10 have all emphasised that such estimates must inform the allocation of resources to an investigation as well as the implementation of any recommendations. The Tärnsjö report also illustrates good practice by the way in which the investigators justify their assessment of this 'worst plausible outcome' . Auch justifications are particularly important if readers are to be convinced by this subjunctive form of reasoning. The report describes how the investigators' finding was reached without access to the Amur-2524's stability data under the relevant load conditions. In spite of this they argue that 'if the timber had not entirely filled the hold but the logs, instead, had been able to move around or shift while floating, the vessels metacentre height could have been decreased by half a metre through the effect of the free water surface alone... (this) could have constituted a serious risk to the vessels stability" [767].

A further strength of the Swedish report is that the analysis extends beyond the immediate and the distal causes of the incident to examine the response to the collision. This analysis extends the initial reconstruction because it does not simply discuss what was done, it also considers what was not done in the aftermath of the incident. The Södertälje Traffic Information Centre was informed immediately but the Marine Rescue Control Centre was not informed until approximately ninety minutes after the collision. Similarly, the analysis goes on to argue that the crew of the Tärnsjö should have remained on-site to assist the Amur-2524 after thad determined that they only sustained minor damage. Instead of illuminating the area and helping the Amur-2524 to break the surrounding ice, the "Tärnsjö's master chose to leave the site right after the accident and before the Amur-2524 began her attempt to reach shallow water" [767]. They had to abandon this plan as it became clear that the vessel was more seriously damaged than had first been appreciated. The investigators acknowledge the inherent difficulty of performing an accurate damage assessment in the aftermath of a collision. The investigators use a form of counterfactual reasoning to argue that the Tärnsjö's haste contributed to Amur-2524's predicament:

> "The Tärnsjö or the tug could have escorted the Amur-2524 back to Hjulsta bridge where they could have obtained help to pump out the bilge and investigate the damage. Even if the vessel had been assisted by the tug during the journey, the Board estimates it would have been more difficult and risky to transfer the crew to the tug if the vessel had started to list or even capsized. The Board considers it remarkable that the Tärnsjö's master chose to proceed immediately following the accident without ensuring that the Amur-2524 was out of danger. Evidently the vessel left the site without having fully ascertained the extent of the damage. He was aware that the Amur-2524 was going to try to reach shallow water, but did not wait to see that she accomplished this." [767].

This counterfactual argument illustrates the importance of looking beyond an initial adverse event to look at the response to an incident. In this case, the consequences to the Amur-2524 could have been far more serious than they actually were. It is interesting to note, however, that this line of analysis is not represented in the conclusions that are drawn by the report. The investigators argue simply that the immediate cause of the collision was that the meeting of the two vessels was "poorly planned" [767]. Contributory factors included insufficient communication between the pilots and ineffective use of the resources available on the bridge. It is also interesting to note the recommendations that were made on the basis of the extensive analysis that was performed by the investigators. The proposed interventions were summarised by a single word: None. This raises a number of questions. In particular, the lack of any recommendations suggests that similar incidents might recur in the future.

Presentation Issues for Causal Analysis

The previous paragraphs in this section have used case studies drawn from the Japanese Marine Accident Inquiry Agency and the Swedish Board of Accident Investigation to illustrate a number of key points about the presentation of any analysis in an incident report. It has been argued that this analysis should consider not only the proximal and distal causes of an incident. The report must also present the findings of any analysis of mitigating actions in the aftermath of an incident. Similarly, it has been argued that incident reports should document the evidence and the analytical processes that support particular findings so that readers can understand why particular conclusions were reached. The following enumeration builds on these observations and presents a number of recommendations for the presentation of analytical material in incident reports. It is important to emphasise that these guidelines are not appropriate for all systems. For example, the informal nature and limited resources of local systems can impose tight constraints both on the analysis that is performed and on the documentation of that analysis. Similarly, it can be difficult to summarise all of the procedures that might be used in order to support the findings of larger-scale investigations. Equally, however, it remains the case that many incident reports present findings that cannot easily be justified from the evidence that is presented in a report [749]. This can lead to skepticism about the benefits of participation and can encourage the creation of 'war stories' that provide alternative accounts of the events that were not revealed in the official report:

1. *describe the process used to derive the findings.*
   It is important that incident reports describe the steps that were taken to identify particular causes. For example, the IMO request information about the 'form' of any investigation that is conducted for more serious classes of incident. The Japanese Marine Accident Inquiry Agency describe how "the Judges make the announcement of the judgement in the court after argument" [393]. Similarly, it is important that the readers of an incident report should be able to identify the validation activities that support particular causal findings. For instance, the Board of Accident Investigation ensures that 'interested parties' have the right to follow an investigation. They can also request "further investigative measures that they deem necessary" [768]. This does not go as far as the right of reply that is granted under sub-regulations 16(3) and 16(4) of the Australian Navigation (Maritime casualty) regulations. These provide 'interested parties' with the opportunity to have their responses published together with the findings of the investigators' analysis. The key point here is that the investigatory process should be apparent to the readers of an incident report. Unless people are provided with this information then they cannot be assured that potential biases will not unduly influence the results of an investigation. The suspicion that findings may be based upon an isolated subjective opinion can be sustained even when investigators have gone to great lengths to ensure the veracity of their findings [199].

2. *document the process used to produce the findings.*
   For many incidents, it may be sufficient simply to outline the processes that were used to identify and validate particular findings. In other situations, however, it is also important that the readers of a report can assess the analytical procedures that were employed. This is important if, for instance, the possible consequences of a future recurrence are judged to be particular severe. Similarly, it is important that readers can critically examine the conduct of any analysis if the proposed remedies impose considerable burdens upon the recipients of a report. Very few organisations that recommend techniques, such as the use of Fault Tree diagrams or ECF analysis, ever publish the documents that are intended to support their findings. This is worrying because it can be difficult to avoid making mistakes with these and similar techniques [529, 424]. In consequence, the readers of a report simply have to trust that investigators and the other members of investigation boards have correctly applied those methods that have been adopted. The role of the outside expert is a particular concern in this respect. Too often, investigation agencies have relied upon the advice of individuals whose reputation and expertise is sufficient validation of their insights. This can be particularly frustrating for other readers with their area of expertise who must struggle to interpret their

findings during subsequent design and redevelopment. Several years ago I wrote a paper entitled 'Why Human Error Analysis Fails to Support Systems Development' [408]. This summarised the problems that I felt when the analysis of human factors analysis in an incident report used technical terms, such as 'crew resource management' or 'high workload', without showing any evidence or reasoning to support these particular findings. This created enormous problems for the safety managers and operational staff who then had to take concrete measures to prevent such failures from recurring in the future. It is insufficient simply to state that these problems occurred without documenting the detailed reasons why such a diagnosis is appropriate.

3. *link evidence to arguments.*
   A particular strength of the Swedish Board of Accident Investigation's approach is that the analysis of an incident is directly linked to an assessment of the evidence that is presented within any reconstruction. The Tärnsjö case study illustrates the way in which the investigation board is content to leave potential inconsistencies, for instance over the time of the collision, where they do not affect their findings. In other situations, the Board actively reject eye-witness statements that are contradicted by other forms of evidence. For instance, observations about the progress of the Tärnsjö are used to infer the speed of the vessel. This is then used to support the main findings of the report that contradict the Pilot's statements. In other situations, the Board clearly state where they could not obtain the evidence that might be necessary to directly support their findings. For instance, the investigators argue that the consequences for the MV Amur-2524 could easily have been far worse even though they could not obtain a detailed stability analysis for the vessel. Such transparency is rare found. Most incident reports simply describe the results of any causal analysis. This leaves the reader to piece together the evidence that supports those findings from the previous reconstruction sections of the report. This creates many problems. Firstly, it increases the burdens on the reader who must also assimilate a complex mass of contextual information that increasingly characterises incident reporting involving high-technology applications. Secondly, there is a danger that readers will fail to identify the evidence that support the investigators' findings. In pathological cases, this evidence may not even have been included within the report [426]. Finally, there is a danger that readers may infer causal relationships that were not intended by incident investigators. In such circumstances it is likely that particular items of evidence, such as witness testimonies, will be given undue significance. This becomes important if that evidence is subsequently challenged or is stronger that the evidence that the investigator used in deriving a particular finding.

4. *justify proportion of blame.*
   As we have seen, most reporting systems employ a proportionate approach in which deliberate violations and illegal acts are handled differently than human error. It can, however, be difficult to distinguish between violations, slips, lapses or mistakes from observations of an incident. Such evidence rarely yields insights into operator intention. Investigators must, therefore, explain why they interpreted observations in a particular manner if readers are to understand the basis for their findings. In other systems, such as the Japanese maritime case study, incident reports result in punitive actions. The importance of justifying those findings is correspondingly greater. This explains the use of a judicial investigation process. Even in 'no blame' systems there is still an obligation to explain why the analysis of an incident focuses on certain causes. For instance, it is important to describe why systems failure might have been considered more significant than operator involvement. 'No-blame' systems also assume additional burdens. By shifting attention away from system operators, they hope to identify the contextual and environmental factors that contribute to incidents and accidents. They must, therefore, look to the higher-level organisational and managerial influences that expose systems and their operators to such 'adverse' conditions. Too often there is a tendency to describe environment and contextual factors as part of a more general 'safety culture' without ever address the detailed resons why such a culture is permitted to exist within a particular organisation. Such approaches not only avoid blaming specific individuals, they also ignore

the importance of managerial responsibility for the consequences of particular incidents.

5. *justify any exclusions.*
   Not only must investigators justify their decision to focus on certain causal factors, it is also important that they document the reasons why other forms of analysis were excluded or eliminated. For instance, the NTSB routinely includes a section entitled 'exclusions' within their analysis of incidents and accidents. In their maritime reports, this typically considers whether adverse weather conditions contributed to any failure. They may also document the investigators' findings that "the vessel's navigation, propulsion, and steering systems had no bearing on the cause of the fire" [606]. Subsequent paragraphs, typically, exclude inadequate training or qualification. They may also consider the potential influence of drugs or alcohol. In some cases this can lead to qualified exclusions. This can be illustrated by a recent investigation into an on-board fire. Tests indicated that the first officer had used marijuana in the weeks before the fire occurred, however, "based on witnesses descriptions of the first officers actions on the bridge during the emergency, no behavioral evidence indicated that he was impaired by drugs at the time of the fire" [606]. Such explicit exclusions help to address the alternative causal hypotheses that readers often form when analysing incident and accident reports. This is particular important because empirical studies have shown that trained accident investigators will often sustain such theories in the face of contradictory evidence. For instance, Woodcock and Smiley describe a study involving incident investigators from the American Society of Safety Engineers. One of the fifteen participants continued to believe the drugs were a factor in an incident even though witness statements contradicted this belief [870].

6. *clearly assess the 'causal asymmetry' of the incident.*
   Chapter 11 introduced the term 'causal asymmetry'. This is used to describe the imbalance that exists between attempts to identify the effects of causes and attempts to identify the causes of given effects. It can be difficult or impossible to determine the precise causes of a given effect. Lack of evidence and the existence of multiple causal paths can prevent investigators from ever finding a single causal hypothesis. This is particularly true of incidents involving human intervention. As we have seen, it can be particularly difficult to identify those factors that influence particular operator behaviours. Previous chapters have described the problems that arise when attempting to identify the ignition source of maritime fires [621]. Two particular effects are at work here. Small-scale, local reporting systems often lack the necessary investigatory resources to conduct the detailed analysis that may be required to unambiguously determine causal sequences. Large-scale national systems are, typically, only able to devote necessary analytical resources to a small percentage of the most serious incidents that are reported. The opposite aspect of causal asymmetry is that investigators must trace a path between the cause that is identified in their findings and the observed effects that are recorded in the aftermath of an incident. These forms of analysis are, typically, constructed using the counterfactual arguments that have been described in Chapters 10 and 11. There are particular problems associated with these forms of argument. In particular, readers often draw 'incorrect' inferences about the potential consequences of any counterfactuals [124]. In practical terms, the problems of causal asymmetry in incident reports are best addressed by validation activities. These can be used to determine whether the potential readers of a report can reconstruct the causal arguments that are presented. These studies can also be used to determine whether readers are convinced by the exclusions that are intended to prune alternative causal hypotheses. These validation activities are described in later sections of this chapter.

7. *present the analysis of catalytic, latent and mitigating factors.*
   Chapters 10 and 11 have described the importance of analysing a range of causal factors. These include both the triggering events that led directly to an incident. They also include the longer term factors that contribute to a failure. We have also emphasised the importance of analysing the mitigating or exacerbating factors that can arise after an initial failure. Many investigators routinely consider these factors but then fail to document them within the resulting report. These omissions can be justified in the summary reports that support the statistical analyses

described in Chapter 15. Even there, however, it can be important for safety managers and regulators to gain a clear understanding of the broader context in which an incident occurs. Too often incident databases simply record the frequency of particular catalytic failures [444]. It, therefore, comes as little surprise when the overall failure rate fails to decline over time. New catalysts replace those that have been addressed by previous recommendations and the latent causes of accidents and incidents remain unresolved. Such potential pitfalls can be avoided if information about latent and mitigating factors is propagated into the reports that document particular investigations.

8. *examine the adequacy of interim recommendations.*
   The analysis in an incident report may also consider whether or not any interim recommendations are sufficient to address those causes that are identified by a more detailed investigation. It might be argued that this form of analysis should be considered together with the recommendations in an incident report. The NTSB follows the practice adopted by a number of similar organisations when it includes this form of assessment in the analytical sections of their incident reports. Investigators categorise each interim recommendation proposed by the Safety Board as either acceptable or unacceptable and as open or closed. For instance, recommendation M-98-126 was drafted following a fire on-board a passenger vessel. This required operators to 'institute a program to verify on a continuing basis that the laundry ventilation systems, including ducts and plenums, remain clean and clear of any combustible material that poses a fire hazard on your vessels' [606]. 21 of the 22 recipients of the recommendation replied to indicate that they had taken measures to comply with the proposed actions. As a result, the Board classified Safety Recommendations M-98-126 as a Closed Acceptable Action.

The previous paragraph described how organisations such as the NTSB often include an assessment of interim recommendations within the analysis section of their incident reports. We have not, however, considered how to include particular recommendations within these documents. The following section, therefore, extend the previous analysis to consider the problems that must be considered when presenting readers with proposal that are intended to avoid future incidents.

13.2.3 Recommendations

Previous sections have described how the executive summaries are usually followed by incident reconstructions. These precede paragraphs of analysis that then support any recommendations that might be made in the aftermath of an incident. Previous sections have also stressed the diverse nature of the publications that are produced in order to publicise information about adverse occurrences and near-miss failures. This has important consequences. In particular, Chapter 12 assumed an investigatory process in which investigators presented their recommendations in a final report to a regulatory organisation. They, in turn, were then assumed to provide feedback on whether or not each recommendation might be accepted for implementation throughout an industry. As we have seen, however, incident reports fulfill a more diverse set of roles. For example, interim reports can be directly targeted at the operators and managers who must implement particular recommendations. Conversely, statistical reports may not be published in any conventional sense. Instead, they often summarise the recommended actions so that investigators can survey previous responses to similar incidents in the future. Investigators must, therefore, tailor the presentation of particular recommendations to both the audience and the intended use of the document in which they appear. This point can be illustrated by the list of proposed outputs that are to be derived from the US Coast Guard's International Maritime Information Safety System (IMISS) [834]:

- Alert Bulletins. These distribute interim recommendations as soon as possible after an incident has occurred. These recommendations may simply focus on the need to increase vigilance in order to detect potential failures that, as yet, cannot be prevented by more direct measures.

- For Your Information Notices. Less critical incidents can result in notices that inform personnel of potential problems. The criticality of the information contained should be clearly distinguished from the more immediate alarm bulletins mentioned above. Any recommendations associated with these notices should be periodically reviewed to ensure that they do not provide stop-gap solutions to the latent conditions for more serious incidents in the future.

- Monthly Safety Bulletins. These documents can present collections of 'lessons learned, safety messages, areas for improvement, precautions and data trends' [834]. They play an important role in informing the wider maritime community of significant safety issues. The Notices and Alerts, mentioned above, can be distributed directly to the safety managers and representatives who must oversee the implementation of particular recommendations. In contrast, Monthly Bulletins often present more general recommendations directly to individual operators in the field.

- Periodic Journal and Magazine Articles and Workshops. Individual incident reports can be analysed by a 'data centre'. They may then write articles that examine the effectiveness of recommendations in terms of safety analyses and trends. These should not only be published through in-house media. They should also be submitted to relevant safety journals or conferences. This provides an important means of obtained external peer review for many of the activities that are conducted within a reporting system.

- Publicly Available Database. Databases can be established to provide the public with access to summary information about previous adverse occurrences and near-miss incidents. Increasingly these are provided over the Internet, this approach will be discussed in the following sections and in Chapter 15. An important benefit of these systems is that they can enable 'interested parties' to follow the actions that have been recommended in the aftermath of previous incidents. If they have been unsuccessful in prevent recurrences then this also can be inferred from these collections.

- Client Work and Research Services. Some incident reporting systems recoup some of the inherent expenditure that is involved in their maintenance by offering a range of additional commercial services. These can include advances search and retrieval tasks over their datasets. It can also include services that support other investigatory bodies. For instance, they may be anxious to identify any previous recommendations that have been made in response to similar incidents in other 'jurisdictions'.

- International initiatives. Incident reports can trigger actions to change international agreements. This can involve the publication of requests to amend the recommendations that have been agreed to by members from many different states. Clearly, the scope and tone of such requests are liable to be quite different from recommendations that are directed at the operators and managers within a particular industry.

- Measures of Success. The data derived about the causes of incidents and the proposed recommendations can be used to derive a number og high-level measures of the success of a reporting system. The problems that can arise from this approach are described in Chapter 15. For now it is sufficient to realise that such activities can identify the need to revise previous recommendations. The publication of this data can, therefore, trigger the publication of additional Alerts and Notices.

Brevity prevents a full discussion of the particular techniques that are used to present incident recommendations in each of these different forms of report. In contrast, the following pages draw upon examples of the most widespread formats, for example Information Notices and Alert Bulletins. These are used to identify general features that also characterise aspects of these more diverse formats. For example, the UK MAIB's Safety Digest is typically only produced twice or three times per year. They are, however, intended to fulfill the same role as the 'monthly bulletin' described in the previous list. These documents present case study incidents and then draw a range of general recommendations that are intended to inform the future actions of many different operators and managers within the maritime industries. This can be illustrated by two recent incidents involving fishing vessels [516]. The first incident began when a wooden fishing vessel touched bottom. This caused damage to the vessel's planking. Shortly after this a bilge alarm in the fish hold went off. The crew discovered that the hold was flooding. They contacted the coastguard and headed for the nearest port. Watertight bulkheads either side of the hold restricted the flooding. The vessel made port safely and was eventually pumped dry by a fire brigade tender. In the second incident, a bilge alarm went off during a severe gale. The crew of the trawler found that their engine room was flooding from a fractured casing of the main engine driven cooling pump. A bilge pump was started from an auxiliary power source. This incident was particularly serious because the damaged casing meant that the main engine had to be slowed in order to reduce the rate of flooding. This compromised the vessel's ability to remain close to a lee shore. The engine had to be shut down when water reached the main engine flywheel. The valve to the sea water inlet was closed to stop any more flooding and the engineers then investigated why the bilge pumping had so little effect. Debris had accumulated inside the bilge line valve body and this prevented it from closing. Once it had been cleaned, the bilge pumping system functioned as intended. A salvage pump from a lifeboat was also used to lower the water level in the engine room so that the main engine could be restarted. These incidents inspired a number of recommendations. It is interesting to note, however, that the MAIB refers to these as lessons. This is an important distinction because the term 'lesson' can to imply the sharing of insight into a particular problem. The term 'recommendation' is reserved for the proposed actions that are identified as part of a formal enquiry into particular accidents. The Safety Digest, in contrast, identifies the following findings from the two case studies described above:

1. "Both cases illustrate the benefits of bilge alarms, functioning bilge pump systems, and watertight bulkheads to limit the severity of a flooding incident.

2. The second incident shows the importance of maintaining a vessel's bilges free of rubbish so it cannot be drawn into the bilge system.

3. Valves on a bilge system must be regularly checked for correct operation." [516]

This case study illustrates the way in which the readers of periodical bulletins can be presented with insights from recent incidents. As can be seen, they typically take the form of reminders. They emphasise the importance of particular safety-related activities, such as the maintenance of bilge lines and valves. They also illustrate the potential consequences of neglecting those activities. The previous example also illustrates the way in which incident reports can be used to emphasise the importance of particular items of safety equipment. This is particularly significant for highly-competitive industries in which their might be an economic temptation to avoid the installation and maintenance of such devices that are not strictly necessary in order to achieve particular production quotas.

It is difficult to capture the diversity of recommendations that appear in incident reports. It can also be difficult to anticipate the particular focus that investigators will take when making particular proposals. Both of these points can be illustrated by the conclusions of a report by the Australian Maritime Safety Authority into livestock mortalities on board the MV Temburong [42]. This incident occurred when the deck generator supplying power to a livestock ventilation system failed. The crew identified that the generator's diesel fuel supply had been contaminated with a heavier grade of intermediate fuel oil. The power for the ventilation system was then transferred to the engine room generators. These subsequently failed. The consequences were a complete loss of electrical power including the shut down of the ventilation system to the livestock spaces. This second failure was traced to water contamination of the engine room generator's diesel fuel supply. The report focuses on the adverse outcome of the incident; over 800 cattle died. Less attention is paid to the potential dangers posed to the crew and other mariners by the loss of power to the vessel. This focus is entirely justifiable given that the investigators associated with a relatively low risk with the second and total power loss. This assessment is never made explicit within the report. The Australian Maritime Safety Authority do not present particular recommendations in this report. In contrast, they enumerate a number of conclusions. This approach is instructive because it typifies incident reporting systems that separate the determination of what happened from the recommendation of potential interventions:

> “14.1 The failure of both the primary and secondary sources of power was due to contamination of the ships fuel supplies: contamination of the fuel for the primary power source was water and the contamination in the case of the secondary power source was heavy fuel oil.
>
> 14.2 The situation was compounded by the failure of the ships communications systems during the incident as a result of an unexplained battery failure.
>
> 14.3 The serious nature of the incident was exacerbated by inadequacies on the part of the ships personnel in their immediate response to the incident and the absence of contemporary ship management practices.
>
> ...
>
> 14.6 On departure the operation and configuration of the ships generating equipment was in accord with the prescribed regulatory requirements. However the manner of operation has raised the issue of whether the Marine Order requires redrafting so as to ensure that the meaning of a secondary power source is more clearly defined.
>
> 14.7 Whilst the fuel contamination situation was avoidable the incident has nevertheless raised the general issue of the ability of livestock vessels to recover from a dead-ship situation. Whilst the Temburong satisfied class and flag rules in this regard the system was shown to be particularly vulnerable when the well being of the livestock is considered.
>
> 14.8 Following the restoration of power the crew performed well in a most difficult situation to remove the dead cattle from the vessel.” [42]

As can be seen the findings from this investigation address a broad range of system failures. They consider problems in the crews’ management, they address the failure of specific subsystems, they also consider regulatory failure and potential problems in the interpretation of Marine orders. The previous quotation does not, however, present particular recommendations. These are left implicit within the findings that are derived from the causal analysis. Sentences such as ‘whilst the the Temburong satisfied class and flag rules in this regard, the system was shown to be particularly vulnerable’ imply that the rules or their interpretation may be inadequate to prevent future incidents. Similarly, the report questions ‘whether the Marine Order requires redrafting so as to ensure that the meaning of a secondary power source is more clearly defined’. This arguably implies that a recommendation be made to examine the sufficiency of the relevant Marine Order. Previous paragraphs have argued that this approach typifies incident reporting systems in which the identification of what happened is separated from the recommendation of remedial actions. This represents an extreme example of the heuristic introduced in Chapter 12 that investigators should leave domain experts to determine *how* to avoid future failures. It can, however, lead to a number of potential pitfalls. In particular, it is essential to ensure that the organisations and individuals who subsequently identify potential recommendations are independent from those who are implicated in a report. Fortunately, the investigators in our case study avoid this potential pitfall by passing their report directly to the Australian Maritime Safety Authority who, in turn, are responsible for identifying means of avoiding similar incidents in the future [43].

The use of language in an incident report can reflect the effect that particular recommendations are intended to have upon their audience. The previous recommendations illustrate the focused use of particular recommendations to reinforce existing safety information. They do not require significant additional expenditure nor do they imply a major safety initiative on behalf of the industry involved. In contrast, very different language is used by the US Coast Guard when a Quality Action Team presents its recommendations into towing vessel incidents on American Waterways [828]. They argued that the incidence of fatalities and injuries can be reduced by a program including prevention measures, the collection and dissemination of lessons learned, improved investigation and data collection techniques and the regular assessment of towing industry performance over time using a fatality rate model. in order to interpret the language used in these recommendations it is necessary to understand the nature of the relationship between the Quality Action Team and the intended recipients of their proposals. The team were anxious to preserve a cordial working relationship between the Coast Guard and the American Waterways Operators Safety Partnership. In consequence, they were anxious to stress 'non-regulatory solutions'. Under prevention, they argued that companies should implement a fall overboard prevention program. In particular, they should:

1. "Formulate and implement fall prevention work procedures consistent with vessel mission. crew complement, and geographic area of operation. Procedures should emphasise teamwork, communication, and safe work practices.

2. Ensure that all crewmembers receive initial and recurrent training in such procedures.

3. Assign responsibility for ensuring compliance with procedures to on-board supervision.

4. Enforce fall overboard prevention policies and consider the use of counseling, recurrent training, and discipline in the enforcement process.

5. Investigate all fall overboard incidents to determine what happened and how such incidents could be prevented in the future.

6. Inform all employees of fall overboard incidents and use lessons learned as part of a recurrent training program.

7. Modify fall prevention procedures as necessary based on investigation of fall overboard incidents." [828]

The detailed nature of these requirements arguably illustrates the way in which the Quality Action Team sought to ensure that the recipients of their report clearly understood the proposed remedies. This contrasts with the previous recommendations from the UK MAIB report that were far less detailed in encouraging companies to follow what can be regarded as existing practices. The previous requirements can be contrasted with further recommendation that were made by the Coast Guard. These acknowledged the diverse circumstances that characterise towing operations. They, therefore, encouraged companies to select 'best practices' from a list that was developed during the drafting of the report. This illustrates the cooperative, 'non-regulatory' tone of the report.

These is an interesting recursive element to the Coast Guard report. Their recommendations include the establishment of an incident reporting systems that would be used to 'publicise the findings and recommendations of the Towing Vessel Crew Fatalities Quality Action Team through the American Waterways Operators Letter, sector committee meetings, regional meetings, and the Interregion Safety Committee' [828]. The recommendations were also applied reflexively to members of the Quality Action Teams own parent organisation. They argued that the Coast Guard should publicise their recommendations through the Marine Safety Newsletter and 'Industry Days'. Brevity prevents a more sustained analysis of a detailed and exhaustive set of recommendations to a complex problem. It is worth noting in passing, however, that this document like many other products of incident report systems breaks many of the guidelines that were presented in Chapter 12. For example, we have described the US Air Force's injunction not to recommend further studies [794]. Instead, investigatory and regulatory bodies should propose actions even if they are are contingent upon on-going studies. In contrast, the Coast Guard's report identifies numerous areas for further analysis. American Waterways Operators and the Coast Guard are urged to find out more about the factors influencing survivability following a fall overboard Incident. They also recommend that the Quality Action Team's report be used to guide a more detailed analysis of fatalities in other segments of the barge and towing industry. Finally, they describe how further studies must be conducted to determine a means of measuring the impact that fall prevention programs, policies, and procedures have upon incident rates. The differences between this Coast Guard report and the heuristics that guide investigatory practice in the US Air Force can, in part, be explained by the differences in the nature of the recommendations that they propose. These differences, in turn, stem from the diverse nature of the reports that are being produced. The Air Force procedures refer to the recommendations that are produced after specific incidents and accidents. Further studies imply a delay that may jeopardise the safety of existing operations. Hence, the focus is on providing a rapid and effective response to particular failures. In contrast, the Coast Guard report presents recommendations that are themselves the product of a study into a collection of incident reports. Given the complex nature of these occurrences and previous problems in reducing the rate of towing incidents it is hardly surprising that some of the recommendations should focus on the need for further analysis.

As mentioned, the form and tone in which recommendations are presented can vary depending on the intended audience and the publication in which it appears. This can be illustrated by the recommendations from a more detailed investigation by the Transportation Safety Board of Canada into a 'bottom contact' incident involving a bulk carrier [785]. This report cannot easily be categorised in terms of the previous enumeration of document types because it does not directly introduce any new recommendations. It does, however, validate the interim measures that were proposed in collaboration with the Canadian Coast Guard . It is important to emphasise, however, that the lack of any new recommendations does not imply that the report does not play a role in the presentation of proposed interventions. In contrast, it provides an important means of disseminating information about those actions that have already been initiated by other organisations. The main cause of the incident was identified as the bridge crew's lack of information about a 'shallow spot' in the navigable channel of the Fraser River. This was attributed to inadequacies in the current system of monitoring navigable channels and producing depth information for vessels in that area. As mentioned, rather than making specific recommendations the investigators described the 'safety actions' that had been taken. The report concludes the Transportation Safety Board's investigation and, therefore, the reader can assume that these actions are considered to be sufficient to ensure that similar incidents will not occur in the future:

> "The Canadian Coast Guard advised that a formal Working Committee, with representatives from the Fraser River Pilots Association, Fraser River Port Authority and Coast Guard has been established and will be meeting quarterly to review channel conditions and status of the channel monitoring and maintenance dredging program. A possibility of modelling the sedimentation process to determine various rates of in-fill associated with forcasted river flow/discharge will be explored." [785]

It is clear from this quotation that the previous recommendations have a relatively limited geographical scope. They focus on the activities of the Fraser River Pilots Association and the Fraser River Port Authority. The local nature of such proposed interventions enables investigators to defer much of the detail about any recommendation. In contrast, previous paragraphs have described how incident investigations can motivate far more general proposals. These recommendations go beyond working practices in particular geographical area to address industry-wide operating procedures throughout the globe. The presentation and form of such recommendations can be very different from those in the interim and final reports that are generated after many incidents. This can be illustrated by the NTSB's recommendations following a series of fire on board passenger ships. The Board issued detailed proposals to the International Council of Cruise Lines and Cruise Line Companies Regarding Fires on Board Passenger Ships. As might be expected, such an action was not taken without considerable resources being allocated to analysing the causes of many previous incidents. Although the letter was drafted in July 2000, the initial incidents that prompted their intervention occurred on board the Universe Explorer in 1996 and the Vistafjord in 1997 [613]. The NTSB issued safety recommendations to the U.S. Coast Guard in April 1997 that automatic smoke alarms should be installed to protect both crew berthing and passenger accommodation areas. These recommendations were, in turn, forwarded by the Coast Guard to the International Council of Cruise Lines and the International Chamber of Shipping. The NTSB's proposals were opposed on the grounds that such systems might generate false alarms and could create crowd control problems. Three further incidents in 1998, 1999 and 2000 prompted the Board to reiterate their recommendations. They requested that Cruise Line Companies 'without delay install automatic local-sounding smoke alarms in crew accommodation areas on company passenger ships so that crews will receive immediate warning of the presence of smoke and will have the maximum available escape time during a fire' [613]. The same recommendation was also made for the installation 'without delay' of the same devices in accommodation areas on company passenger ship. The International Council of Cruise Lines were requested to withdraw their 'opposition to the amendment of the Safety of Life at Sea Convention chapter II-2 to require automatic local-sounding smoke alarms in crew accommodation spaces on board passenger ships and support a full discussion of the technical issues and any further U.S. Coast Guard actions on this matter before the IMO'. The same request was reiterated for their opposition to smoke alarms in passenger accommodation spaces. The recommendations concluded with an appeal that the council should support a 'full discussion of the technical issues involved and any further U.S. Coast Guard actions on this matter before the IMO' [613]. It is important not to overlook the colour or tone of the language used in this document. The note of exasperation or frustration reflects the Board's concern over this issue. Only time will tell whether this choice of language was an appropriate means of ensuring international agreement over these particular recommendations.

The analytical procedures that have been described in the previous section have a strong influence both on the recommendations that are derived from an investigation and also on the manner in which any proposed interventions are described in a subsequent report. For instance, the Swedish Board of Accident Investigation's 'no blame' approach is deliberately intended to help them issue recommendations as soon as possible after an incident has occurred; "...with the ongoing technical development it is also necessary to quickly get knowledge of shortcomings in an operation that can cause an accident or contribute to one" [768]. In more judicial systems, the focus is on identifying recommendations in a more deliberate fashion given that their findings can have a profound implication upon the livelihoods of the individuals who are affected [393]. The majority of incident reporting systems can be interpreted to lie between these two extremes. Proportionate blame approaches, such as that adopted by the US NTSB, will often issue interim recommendations that are then supplemented by any subsequent findings from a secondary investigation. The previous section has also described how reports assess these initial recommendations in terms of whether or not they were acceptable and, if so, can their implementation be declared closed. This two stage process creates considerable problems for those who must publish and disseminate information about potential recommendations. They must firstly ensure that all of the intended recipients receive copies of initial advisories. It is important that these individuals and organisations both understand the intended response to such recommendations and that they have the adequate resources to satisfy particular requirements. Regulators must then, typically, ensure that the same recipients of the initial advisories are then provided with the updated recommendations that may be made in any particular incident report. Of course, these may also 'countermand' or replace recommendations that were made as a result of incident reports that were issued many months or years before the present investigation.

Such dissemination activities require considerable logistical support. Investigatory organisations have, therefore, developed a range of databases to monitor and support the dissemination of incident recommendations. For example, Chapter 12 describe some of the applications that support US Army reporting systems. Other organisations have developed billboards that help the intended recipients of a recommendation to monitor amendments and revisions. Such resources are important when logistical barriers prevent regulators and investigators from guaranteeing the delivery of proposed interventions. For example, it can be difficult for the crews of many different merchant vessels to continually monitor the findings of incident reports that cover all of their possible ports of call. Later sections will focus on these dissemination problems in more detail. For now it is sufficient to stress that many of these systems fulfill a dual purpose. Not only do they help monitor the distribution and receipt of particular recommendations, they can also be used to monitor their implementation. These databases, typically, restrict access to such information. In contrast, the New Zealand Transport Accident Investigation Commission maintains a World Wide Web list of information about marine safety recommendations [630]. They recognise that "not all safety recommendations and responses are published in the Commission's Occurrence Reports because at the time of printing some recommendations may have been incomplete or some responses may not have been available". The bulletin-board is accessed via a list of previous incidents. For example, the user would select a link labelled 'Report 00-211, harbour tug Waka Kume, loss of control, Auckland Harbour, 19 November 2000' in order to view any revisions to the recommendations that were made in that particular report. An alternative approach is to publish revised recommendations in terms of particular safety-related themes, such as transfers between ships and helicopters. This approach is similar to that adopted by the Australian Maritime Safety Authority [43]. Assuming that the reader had selected the New Zealand link, cited above, they would be presented with information about the recommendations that were made in the initial report. The Commission recommended to the manager of marine services for the Ports of Auckland Limited that he revise their tug operator training manual to include detailed information about engine and control system failures. They also requested that he introduce regular, documented peer reviews to ensure that all tug operators complied with the relevant safety regulations. These reviews should also assess the training that the Ports of Auckland Limited provides for its skippers. They should be performed by independent experts in the operation of similar tugs. The New Zealand Transport Accident Investigation Commission's recommendation summary goes on to describe the Manager's response to these proposals. This response illustrate the way in which bulletin boards can be used to inform 'interested parties' that particular recommendations have been accepted and are being implemented. The Manager confirmed that:

1. "Updates to the tug operator training manuals shall be carried out as follows: engine control system failures, this section shall be expanded; response to engine control system failures, this section shall be expanded; handling the tug with one azimuthing unit, this task shall be given greater emphasis; introduce a system of peer review, this is being done.
2. The requirement to review Ports of Auckland Ltd training skippers by independent experts is not practical. The trainer we originally used was from Canada no other trainer exists in NZ. Many of our staff have gained good skill levels with these vessels and will be adequate for use in peer reviews." [629]

Unfortunately, further problems complicate the presentation of those recommendations that are made in aftermath of an adverse occurrence or near-miss incident. In particular, incident reporting systems must not simply present the findings that are derived from isolated near-miss incidents or adverse occurrences. They must also consider interventions that address common features amongst a number of previous failures. The problems of presenting these findings from multiple incidents are assessed in later sections. In contrast, the following paragraphs focus more narrowly on guidelines for the presentation of recommendations that relate to single incidents.

Presentation Issues for Recommendations

As we have seen, previous incident reports have been weakened by factual omissions. This can prevent readers from gaining a good impression of the events leading to an adverse occurrence or near-miss incident. Similarly, if parts of an analysis are omitted then it can be difficult to follow the conclusions that are derived from an incident investigation. The level of quality control that can be observed by reading large numbers of incident reports is, arguably, higher in the preliminary sections of these documents than it is in the concluding paragraphs that, typically, list any potential recommendations. One means of validating this claim is by looking at the litigation that follows from many incidents and accidents. These proceedings, typically, focus on the proposed interventions in the aftermath of an adverse occurrence. Far fewer proceedings are initiated to question the evidence that is put forward in an incident report. Such differences can also be explained in terms of the consequences that recommendations have upon the future operation of safety-critical systems. Incident reconstructions are likely to be less contentious than any proposed interventions because they need not carry with them the costs that are associated with implementing any subsequent changes. It makes little practical difference whether one explains the focus on recommendations as being due to flaws in their presentation or to natural concerns over the cost implications their implementation. In either case, it is clearly important that analysts devote considerable time and attention to ensuring that their findings are presented in a clear and coherent manner. As we shall see, this can involve the publication of preliminary or interim reports to solicit comments of proposed interventions. It can also involve more restricted forms of pre-publication or consultation during which analysts validate their recommendations before disseminating them more widely. Such techniques can be used to ensure that the presentation of recommendations satisfies a number of detailed requirements:

1. *Identify the recipients and draft recommendations accordingly.*
   Chapter 12 has already described how recommendations should consider what must be achieved but not necessarily how those goals will be implemented. It has also been argued that the recipients of a recommendation must have the necessary resources to achieve particular objectives by the dates that are specified in an incident report. Unfortunately, many incident reports fail to follow these guidelines. For example, the following excerpt is taken from a report issued by the Hong Kong Marine Accident Investigation Section:

   > "Non compliance of the safety guidelines for the transport of motor vehicles/cycles is considered the principle cause of this accident. Had the guidelines been followed, there would have been no uncleaned fuel tanks of motorcycles containing residue of volatile hydrocarbon based substance and residual fuel in the engine assemblies, and as a result no accumulation of hydrocarbon vapour in the container." [364] http://www.info.gov.hk/mardep/dept/mai/elmi399.htm

   This illustrates how recommendations may often be implicit within the findings of an incident report: operating companies must follow the safety guidelines for the transport of motor vehicles. There are a number of problems with this implicit approach. Firstly, reminders to operators that they must 'try harder' offer few guarantees about the future safety of an application process [409]. Secondly, the implicit nature of such recommendations can lead to a range of different interpretations of the proposed remedies. Readers might infer that the harbour authorities ought to initiate more frequent checks to ensure compliance with these regulations. Alternatively, improved training might be offered to the particular working groups that were involved in loading this vessel. In an ideal world, such alternative interpretations might encourage readers to initiate a range of interventions. In practice, however, there is a danger that any ambiguity can encourage individuals and organisations to pass-on responsibility for implementing particular recommendations.

2. *consider the problems of non-compliance or opposition.*
   The most observant readers of this work will have noticed a particular trend that is apparent in the response to many of the recommendations that have been proposed by the NTSB. In Chapter 9 we looked at the recommendations that were issued in response to a succession of gas explosions. The Office of Pipeline Safety in the Department of Transportation questioned NTSB initiatives to install Excess Flow Valves [588]. Similarly, previous sections in this chapter have described the reluctance of the International Council of Cruise Lines and the International Chamber of Shipping to introduce automatic smoke alarms into crew berthing and passenger accommodation areas [613]. One explanation for the opposition that is often provoked by NTSB recommendations is that their investigators focus primarily on the safety issues rather than the cost implications of implementing particular proposals. As we have seen, other incident reporting systems take a more cautious approach when they 'target the doable'. The reluctance to implement particular recommendations, arguably, provides a further illustration of the relationship that exists between corporations and federal agencies in the United States. In either case, such opposition illustrates the importance of explicitly considering what to do when the recipients of a report object to particular recommendations. As we have seen, some reporting agencies simply document the findings of an investigation without providing the reader with any idea of whether or not they were actually implemented. Any objections that block the implementation of a recommendation are then, typically, only revealed in subsequent reports that document the recurrence of similar incidents. Of course, this approach cannot be used to identify situations in which objections to a recommendation have not contributed to

subsequent incidents. This approach creates considerable problems both for incident investigators and safety managers. It can be difficult for them to recreate the different arguments that support and oppose particular recommendations. As a reader of incident reports, it can often be frustrating to see investigatory agencies repeatedly advocate the same interventions without any explanation of why their recommendations are consistently not being implemented. Other reporting organisations have avoided these problems by publicising any dialogues that take place with 'interested parties'. This can be illustrated by the way in which the IMO have collated the arguments for and against particular recommendations in the aftermath of Erica oil pollution incident off the French coast. The IMO collated and published responses from the IMO Maritime Safety Committee and their Maritime Environment Protection Agency. These addressed recommendations and proposals made by the International Association of Classification Societies and a resolution by the European parliament. Their views on changes to condition assessment schemes provoked further response from the International Association of Ports and Harbours. Amendments were proposed by maritime organisations in Belgium, Brazil, the Bahamas, Germany Greece, Japan etc. The IMO Maritime Environment Protection Agency then commented on these national submissions. Further statements were made by the International Chamber of Shipping and the International Association of Independent Tanker Owners. Such multi-national responses represent an extreme example. In most incident reports, it is possible to represent different attitudes to particular findings within a final report. This enables investigators to document the reasons why particular groups might oppose the implementation of any potential recommendations. This explicit statement of objections is very important. There have been many instances in which operating companies have denied objecting to recommendations that might have prevented incidents and accidents.

3. *define conformance and validation criteria.*
Like most guidelines, it is possible to identify a number of problems that must be addressed when attempting to use the items in this list as a means of informing the presentation of recommendations that are derived from incident reports. For instance, we have argued that investigators should, if possible, avoid trying to specify the precise mechanisms and procedures that might be used to satisfy a particular finding. This create problems because regulatory authorities, safety managers and operators must determine appropriate means of meeting particular recommendations. There is an obvious danger that incident investigators might conclude that any failure to prevent subsequent incidents indicates a failure to satisfactorily implement their recommendations rather than concluding that their previous recommendations were inadequate. It is, therefore, important that if a report documents what must be done rather than how then the report must also specify conformance criteria that can be used by operators and regulators to determine whether particular mechanisms have satisfied particular high-level objectives. Or put another way, without such criteria investigators may continue to blame inadequate implementation of previous recommendations rather than investigating whether those recommendations were adequate in the first place. Establish conformance is non-trivial. We might like to specify that any proposed changes will make it extremely unlikelihood that an incident will recur. As we have seen, however, many safety-critical incidents have an extremely low probability. If we consider an incident that occurs once in every 100,000 hours of operation, we may have to wait a considerable period of time before we can have any assurance that proposed changes have reduced this frequency. Basic probability theory suggests that even if we observe a system for 99,999 hours there is no guarantee that an incident will not occur twice in the remaining hour. These issues will be discussed in greater detail in Chapter 15. For now, the UK Coastguard Agency's Marine Pollution Control Unit's report into the Sea Empress Incident illustrates some of the points mentioned in this paragraph [793]. The report does not focus on the causes that contributed to the grounding of the vessel and subsequent release of 72,000 tonnes of crude oil into Milford Haven. Instead, it analyses the environmental response and makes a number of recommendations that are intended to improve the 'clean-up' operation after future incidents. The Sea Empress was grounded in 1996, since that time there have been no comparable incidents that might be used to judge whether or not the

recommendations have had their intended effect. All that can be done is for the relevant agencies to test their preparedness in simulated exercises that are intended to demonstrate that they meet the recommendations outlines in the Coast Guard report. The nature of these exercises can vary significantly from one local authority to another. This guideline, therefore, argues that incident reports should document acceptance criteria that can be used to determine whether or not particular mechanisms have actually satisfied a recommendation.

4. *Link the recommendations to the analysis.*
The previous section has argued that incident reports must develop explicit links between the evidence that supports a reconstruction and the analysis that leads to certain causal findings. This is intended to reassure readers that those findings are grounded in the evidence that is obtained in the aftermath of an incident. Similarly, it is important that an incident report documents the relationship between particular recommendations and the products of a causal analysis. Such links create a clear path between the events that contributed to previous failures and the proposed interventions that are intended to prevent future recurrences. The importance of these connections can be illustrated by two recommendations from a report that was issued by the UK MAIB [515]. The incident began when the watchkeeping motorman on a roll-on, roll-off cargo vessel started to clean the top of an electrical cabinet. To gain access, he stood on top of a pipe and lent over a fuel oil booster pump. As he did this, he inadvertently activated the emergency stop button on the pump. The pump stopped and this caused a fall in the fuel pressure for the main engine. This, in turn, caused the main engine to stop resulting in a 'black out' when the generator breaker on the main switchboard tripped. The particular details of this incident are less significant in the context of this section than the recommendations that were identified. The first of these stressed that all emergency stop buttons should be fitted with bright protective covers that should alert operators to their presence and function. The relationship between this recommendation and the details of the incident is relatively clear given the details that were provided in the reconstruction. In contrast, a further recommendation focussed on the "need for an active Safety Committee" even though the report tells us nothing at all about the activities of the safety committee of the vessel involved in the incident. A recommendation that "safety committees should examine all new regulations, operating requirements and the implications of new equipment to see whether any affect the risk profile" seems almost incidental to the occurrence being reported. Such proposals require further justification if readers are to be convinced that they might play a significant role in prevent future incidents [124]. Conversely, it is also important to justify a decision not to derive any recommendations from an incident. This can be illustrated by the conclusion to the Swedish Board of Accident Investigation's report into the collision between the MT Tärnsjö and the MV Amur-2524 mentioned in previous paragraphs [767]. This tersely summarises the proposed interventions as follows: "4. Recommendations. None". When investigators propose particular interventions, they must provide safety managers with a justification that will motivate them to implement any recommendations. Conversely, if investigators decide not to propose any potential interventions then they must provide the authorities that commission the report or which regulate the reporting system with a justification for their decision that no recommendations need be drawn.

5. *explain why recommendations have been rejected or modified.*
The previous example provides an extreme case in which no recommendations were identified in the aftermath of what might have been a very serious incident. More generally, it is necessary to explain why particular recommendations were rejected rather than justifying a decision to entirely reject making any recommendation at all. This is important because readers often form particular hypotheses about potential interventions that might have avoided an incident. If investigators propose alternative remedies then there is a danger that operators and safety managers will choose to follow their own intuitions rather than the proposed alternative. As we have seen, this can have the knock-on effect of complicating any attempts to validate recommendations unless significant resources are used to assess the degree of non-compliance within an industry. The US Coast Guard avoid this problem by a relatively sophisticated process that

also helps to validate particular recommendations. The initial investigation restricts itself to an examination of the facts that are known about an incident. A subsequent section of 'Findings' are then used to present the results of the analysis. A preliminary report is then presented together with a letter containing proposed recommendations is then sent to the Coast Guard Commandant. They review each recommendation, typically, with the Commander of the district in which the incident occurred. They then draft a letter that is printed at the start of each report. Their response lists each recommendation and states whether or not they concur with the proposal. If they do not concur with the specific recommendation, they may agree with the intent and modify the proposal. It is difficult to find any examples of investigator's recommendations that have been rejected. However, this style of report does present detailed arguments to explain why some recommendation are considered inappropriate. This, in turn, can help to justify the alternative recommendations that are proposed by the Commandant and his staff [831]. The Transportation Safety Board of Canada report into the Fraser river incident, mentioned in previous paragraphs [785], provides an example of an alternative approach. In this instance, rather than having an external auditor validate the recommendations made by the investigators, the investigators validate the interim actions that have been taken by local and regional officers in the aftermath of the incident. As mentioned previously, the report into this incident describes the investigation as closed without proposing any new recommendations. The investigators do not modify or revise the interim recommendations presumably because they are considered to be sufficient to prevent any recurrence of a similar incident.

6. *Distinguish between recommendations and lessons learned.*
This section began by listing the publications that are used to report on information submitted to the US Coast Guard's International Maritime Safety System. These range from formal alert bulletins that describe interim recommendations through to the informal information notices that summarise more general problems and which provide background information about potential hazards. This distinction between different forms of publication provides readers with important cues about the relative importance of particular recommendations. Alert bulletins, typically, describe particular actions that should be allocated an extremely high priority within daily working practices. In contrast, information notices may provide general advice that often simply describes 'best practice' without requiring the rapid implementation of particular procedures. Such differences are apparent both in the medium of publication and in the style of prose that is used to convey these different types of information. For instance, the difference between formal alerts and information notices is intended to reflect a distinction between incident recommendations and more general 'lessons learned'. The Coast Guard describe the style and format of lessons in the following terms: "Lessons learned: the information presented through the links below is mostly anecdotal and primarily intended for those who work on the many vessels that navigate our oceans and waterways" [832]. In contrast, incident recommendations follow the more structured format suggested by the previous guidelines in this section.

This section has presented a number of high-level guidelines that are intended to help investigators draft reports into adverse incidents and near-miss occurrences. The focus has been on the more formal style of report that are typically produced in the aftermath of failures that carry relatively severe potential consequences. This is justified by the observation that many of these reports have been flawed by omissions and inconsistencies [426]. It is also possible to selectively use subsets of these guidelines to inform the presentation of less formal reports. In both cases, it is important that investigators consider means of validating the documents that they produce. The guidelines presented in this section are no more than heuristics or rules of thumb. They are the product of experience in generating incident reports. They may not, therefore, provide appropriate guidance for the ast range of incident reporting systems that are currently being implemented. It is, therefore, often necessary to conduct further studies to increase confidence that incident reports provide their readers with all necessary information in a format that supports the operation and maintenance of safety-critical applications.

13.3 Quality Assurance

This section presents a range of techniques that can be used to support both the validation and the verification of incident reports. Verification establishes that a document meets a certain number of technical requirements. These include the need to ensure that a report does not contain inconsistent information about the course of events leading to an incident. They also include the requirement to ensure that any recommendations do not rely upon contradictory lines of analysis. In contrast, validation techniques can be used to establish that a document actually satisfies a range of user requirements. For instance, it is important to determine whether or not the potential readers of an incident report gain a clear understanding of any proposed recommendations. Similarly, it should be possible to demonstrate that the recipients of a report will have confidence both in the reconstruction of events and in the analysis that is derived from any reconstruction. It can be argued that properties such as consistent and a lack of contradiction are basic user requirements and hence that verification techniques form a subset of the more general validation methods. We retain a clear distinction between these different approaches by assuming that verification techniques often yield insights into a document without the direct involvement of the intended recipients. In contrast, validation techniques involve user testing and the observation of readers using incident reports to support particular activities.

13.3.1 Verification

It is possible to identify a range of properties that investigators might require of an incident report. A small subset of these can be summarised as follows:

- *consistency.* It is important that an incident report should 'get the facts correct'. This implies that the timing of events should be reported consistently throughout a document. If there is genuine uncertainty over the timing of a particular action then this should be made explicit. This form of temporal coherence should be supported by location coherence. The position of key individuals and systems should be consistently reported for any particular moment during an incident. In other words, people should not appear to be in two places at once. Similarly, technical details such as the serial number, type and operating characteristics of particular devices should not change throughout a report unless such changes are explained in the supporting prose. It might seem unnecessary to state these requirements. However, a number of previous studies have document numerous violations of these apparently simple requirements. One of the best known instances involves an incident report in which inconsistent timings were given throughout the document because the emergency services disagreed over who reached the scene first. This disagreement was not made explicit in the document and the reader was offered no explanation as to why key events were given different times in different sections of the report [502].

- *lack of contradiction.* Contradictions can be seen as a particular form of inconsistency. For instance, mathematical proofs of consistency can be used to identify potential contradictions. In incident reports, they occur when the same document assets that some fact A is both true and not true. It is relatively rare to find factual contradictions within an incident report. It is more common to find that an event A is reported to occur at a range of different times rather than an assertion that A did *not* occur at a particular time. Contradictions are, however, more often found in the arguments that support particular findings in an incident report. They occur when the same evidence is used both to support and to weaken particular lines of analysis. As we shall see, the same events can be interpreted within the same report as evidence that operators were following standard operating procedures and yet were disregarding particular safety requirements. It is important that investigators identify such situations not because they are in some sense 'incorrect' but because they often require further explanation in order to convince readers that they do not reflect a deeper weakness within the interpretation and analysis of an incident.

- *limited use of rhetorical devices.* Chapter 11 has described a range of biases that affect the interpretation of information about safety-critical incidents: author bias; confidence bias; hindsight bias; judgement bias; political bias; sponsor bias; professional bias; recognition bias; confirmation bias; frequency bias; recency bias; weapon bias etc. These often stem from the external pressures that social and political groups can place upon investigators. There is a danger that these pressures can encourage individuals to support lines of reasoning that may not be directly supported by the available evidence. There is also a danger that this will result in reports that use rhetorical devices, which help to persuade readers that alternative hypotheses should be discounted. As we shall see, it is impossible within any prose document to entirely avoid the use of rhetorical devices. It is, however, important that investigators are aware of the potential impact of these techniques. It is also possible to employ structured reading techniques to ensure that these devices are not used in a way that might mislead the intended audience about the course or causes of adverse occurrences and near-miss incidents.

- *structural simplicity.* There are a number of structural problems that complicate the task of drafting an incident report. The standard formats mentioned in previous sections require that reconstructions and summaries of the events leading to an incident are presented before any sections that analyse the causes of an incident. Dozens of pages can, therefore, separate the presentation of evidence from the arguments that use particular facts about the course of events. Similarly, recommendations may be presented many pages after the lines of analysis that support them. Many investigators have addressed this problem by reiterating factual information within the analysis section. Analytical arguments may also be repeated before the presentation of any recommendations. This creates further problems because subsequent editing of either section can introduce the inconsistencies and contradictions mentioned above. In other reports, key facts can be omited from a reconstruction and may only be presented before particular recommendations. Previous sections have referred to such practices as the 'Perry Mason' or 'Agatha Christie' approach to incident reporting. Readers cannot form a coherent picture of an incident until they have read the closing pages of the report.

This is a partial list. Investigators can identify a range of other requirements that should be satisfied by incident reports. For instance, they might wish to guarantee that sufficient warrants are provided to support the evidence that is presented in these documents. Warrants describe the backing or source that supports particular items of information. Brevity prevents a complete analysis of these diverse properties. In contrast, the following pages briefly introduce a number of techniques that can be used to verify particular properties in incident reports. As mentioned, they can be distinguished from validation techniques because they can typically be performed by an analysis of the document prior to publication and they do not require direct access to the intended recipients of a particular report.

Analysis of Rhetorical Devices

Rhetorical devices, or tropes, represent common and traditional techniques of style and arrangement that can be used in prose to achieve a number of particular effects. These effects include emphasis, association, clarification and focus. They also involve the physical organisation of text through transition, disposition and arrangement. A further class of tropes deals with variety and decoration. It is difficult to write effective prose without making use of these different devices. It is also important to understand the particular effects that these devices can have upon the readers of an incident report. In particular, investigators must recognise when their use of particular rhetorical devices might have an unwanted or unwarranted impact upon their audience.

There are many good textbooks that provide an overview of rhetorical techniques. Some, such as the primer by Corbett and Connors, also include progymnasta or classical composition exercises that are intended to help writers use particular devices [183]. It is ironic that many of these techniques, that are traditionally intended to persuade others of your own opinion, might be abused in the context of incident reporting. They provide case studies in the use of language to create an effect that does not stem solely from the information that is contained in the sentence. Brevity prevents a complete introduction to all of these techniques. In contrast, the following list briefly summarises a selection of the most common tropes and provides examples of their use within a single incident report. These illustrations come from a report that was published by the Transportation Safety Board of Canada [784]. In this incident, the Navimar V, a pilot boat, came alongside a bulk carrier, the Navios Minerva. As she did so, she overtook a wave generated by the carrier and pitched onto its crest. The Navimar V then surged into the trough of the next wave and plunged into the sea. The submerged bow slowed the pilot boat but her momentum continued to pitch the vessel until she turned over. The incident resulted in serious injuries to one crew member and minor injuries to three more people, including two pilots who were using the accommodation ladder to board the Navimar V. The subsequent report into this incident can be used to illustrate the role that particular tropes or rhetorical devices can have upon the reader of an incident report:

- *Amplification*
  This technique involves the restatement of an idea or argument. It often also involves the introduction of additional details. For example, the report into the capsize of the Navimar V includes the following findings:

  > "2. Under international regulations, the vessel was required to use a pilot ladder to transfer pilots.
  >
  > 3. Instructing Marine Communications and Traffic Services traffic regulating officers to tell foreign crews to use the accommodation ladder is a request made by the pilots but is contrary to international regulations." [784]

  The first finding notes that international requirements require the use of a pilot ladder. The second finding amplifies this observation by noting that the request to use an accommodation ladder is also forbidden. This technique can have the effect of drawing the reader's attention to a particular concept or idea. The amplification not only introduces new facts but it also supports and reiterates the arguments that are introduced in previous sentences. This technique can create problems when the amplification of particular aspects of a previous assertion can detract from other arguments or items of information. The combined effect of these confirmatory statements can have a significant impact upon the reader of an incident report. It is, therefore, also important to provide sufficient evidence to establish the credibility of both an initial assertion and the subsequent amplification. For instance, the reader of the Navimar V report is referred to Chapter 5, Rule 17 of the 1974 International Convention for the Safety of Life at Sea (SOLAS). This stipulates that pilot ladders must be used for pilotage service.

- *Anaphora*
  This technique uses repetition at the beginning of successive phrases, clauses or sentences. It can create an impression of climax in which the repetition leads to a particularly important insight or conclusion. The Navimar V report uses this technique in the paragraphs of analysis that immediately precede the investigator's conclusions:

  > "Because the bulk carrier was moving at a speed through the water of about 10 knots, the waves she generated were probably unusually large. During the transfer, the pilot boat was manuevred onto the crest of one of those waves which was just forward of the accommodation ladder platform. Because the speed of the wave was less than the speed of the pilot boat, the 'NAVIMAR V' accelerated into the trough towards the next wave. Because the accommodation ladder was on the vessel's quarter, the bottom platform was not resting onto the vessel's side and the pilot boat's bow had to be rested on the vessel's side to line up the after deck under the accommodation ladder".[784]

  This example illustrates the successive use of the phrase 'Because the...' to build up a causal explanation of one aspect of the incident. It is reminiscent of the phrases that can be derived using the Why-Because Analysis (WBA) described in Chapter 10, although this technique was not used in this instance. The rhetorical device creates the impression of 'building up a case'.

The investigator uses each successive sentence to 'stack up' the evidence in a manner that supports the analysis. It is important to emphasise that such techniques are not of themselves either 'good' or 'bad'. Rhetorical devices can be used to convince us of well-justified conclusions or to support half-baked theories. It is important, however, to be sensitive to the effects that such techniques might have on the readers of an incident report. For instance, the previous citation can be interpreted to provide readers with a clear summary of the arguments that support the investigators' conclusions. It can also be interpreted in a more negative light. The repetition of such phrases may create an impression of certainty about the causes of an incident that might not be justified by the evidence. This particular rhetorical device leaves little room for suggesting alternate hypotheses as this particular argument is being presented.

- *Antithesis*
  This uses juxtaposition to contrasts two ideas or concepts. This can be illustrated by the use of the term 'rather than' in the following sentence: "The custom on the St. Lawrence River for a number of decades has been for pilots to use the accommodation ladder rather than the pilot ladder, unless exceptional conditions require a departure from that practice" [784]. Here the practice of using the pilot ladder is being juxtaposed with the use of the accommodation ladder. This technique is important because readers may make a number of additional inferences based upon such constructions. In this context, it is tempting to infer that the prevailing custom on the St Lawrence is some form of violation. We know that an incident has occurred and the juxtaposition of existing practice with another procedure, such as the use of the pilot ladder, suggests ways in which those alternative practices might have avoided the failure. It can be argued that this analysis makes too many assumptions about the inferences that readers might draw from such an antithetical device. It is important to remember, however, that these additional inferences are very similar to those that been identified by Byrne and Tasso's experimental studies of counterfactual reasoning [124].

- *Asyndeton*
  This device omits conjunctions between words, phrases and clauses. This technique creates an impression of 'unpremeditated multiplicity' [307]. The document creator can think of so many elements in the list that they hardly have time to introduce explicit conjunctions. This is illustrated by an enumeration describing the commitment of Transport Canada: "to the review and approval of construction plans, stability data, and the subsequent inspection of proposed or existing vessels, to ensure compliance with applicable regulatory requirements." [784] The omission of the final conjunction implies that the list is incomplete. There is also a sense in which this technique also builds to a particular conclusion. The final term 'to ensure compliance with applicable regulatory requirements' may be the most important. Asyndeton creates precise and concise summaries. It, therefore, offers considerable stylistic benefits to more formal incident reports that can otherwise appear to be verbose examinations of complex failures. There are, however, dangers. The implied omission of closing conjuncts, as in the previous example, can lead to uncertainty if the reader is unsure of what other information might have been omitted from a list.

- *Conduplicato*
  This technique relies upon the repetition of key words or phrases at, or very near the beginning, of subsequent sentences. This can be illustrated by the following quotation in which the investigators stress the effects of "modifications" on the trimming characteristics of the Navimar V.

  > "The series of modifications carried out in 1993, 1996, and in the spring of 1997, prior to this occurrence, were collectively unsuccessful in eliminating the perceived shortcomings in the vessel's dynamic trimming characteristics. Further modifications, after refloating the vessel in 1998, were made to reduce once again the detrimental after trimming characteristics." [784]

Conduplicato provides a focusing device because writers can use it to emphasise key features in preceding sentences. This helps to ensure that readers notice key concepts or ideas that may have overlooked when they read the initial sentence. The previous example also illustrates the way in which particular passages can simultaneously exploit several different rhetorical techniques. Not only is conduplicato used to emphasise the importance of "modifications" to the vessel. The term "trimming characteristics" is also emphasised by its use at the end of both sentences in the previous example. This technique focuses the readers' attention on the consequences of those modifications. Most investigators draft incident reports without ever being aware that they are exploiting such rhetorical devices. However, some of the material in this section is widely taught within courses on technical writing and composition. The intention behind these courses is to make writers more aware of the techniques that they can exploit in their documents. Previous paragraphs have stressed the dangers that can stem from the (ab)use of particular rhetorical techniques. They have also described how the inadvertent use of some devices can encourage readers to form hypotheses that were not intended by those who drafted an incident report. It is also important to acknowledge that some familiarity with the application of tropes might improve the prose that it often used to present these documents.

- *Diazeugma*
  This rhetorical device involves sentences that are constructed using a single subject and multiple verbs. It is frequently used in the reconstruction sections of incident reports and can help to describe a number of consecutive actions. Diazeugmas can provide an impression of rapid change over time. For instance, the Navimar V report describes the master's actions in extricating himself from the capsized vessel: "he saw light which he thought was coming from the surface and swam in that direction, but found himself in the engine compartment" [784]. As can be seen, this rhetorical device captures a sense of urgency in addition to the temporal information that may be implicit within such structures. There is, however, a danger that diazeugma can provide misleading information if the actions occur over a prolonged period of time. There is also a concern that the implied sequences may divert attention from intervening events. This problem can arise from the use of a single subject throughout the sentence. The actions of other subjects may be postponed to subsequent sentences even though they may have been interleaved with those of the initial subject. For instance, the passage cited above continues as follows:

  > "In the meantime the deck-hand, who was wearing a flotation device, surfaced near the hull. One of the two relief pilots hurried to the bulk carrier's bridge to inform the bridge team of the situation. At 0012, he reported the accident to the Quebec MCTS centre" [784].

  Chapter 9 describes the problems that can arise when readers must reconstruct partial timelines from such prose descriptions.

- *Expletive*
  This technique can be used to emphasise particular concepts by interrupting normal syntax. Examples include the use of 'in fact' or 'indeed'. The following excerpt drawn from the Navimar V report uses the expletive 'moreover' as a preamble to form of amplification. This again illustrates the way in which tropes can be combined to particular effect:

  > "More recent pilot boats generally have a larger embarkation area forward of the wheel-house than aft of it, which makes it easier for the master to observe the transfer manuevre. Moreover, today's pilot boats operate at higher speeds during pilot transfers. Consequently, specific attention must be paid to their dynamic longitudinal trimming characteristics in the design stage, to ensure a safe operation throughout the vessel's displacement, transition, and full-planing modes." [784]

  Such techniques can be used in a variety of different ways. In this example, the expletive is simply used to focus attention on an additional factors that increases the importance of considering the trim characteristics of pilot vessels. There are other instances in which reports have used expletives to cat doubt upon particular aspects of a testimony. For instance, the NTSB report into the loss of a clamming vessel contains the following sentence: "Mr Rubin also testified that he 'absolutely' registered his emergency position indication radio beacon with the National Oceanic and Atmospheric Administration" [831]. In this case the witnesses own rhetorical use of the expletive 'absolutely' is deliberately cited as a precursor to subsequent arguments questioning the truth of their statement. The witnesses own emphasis, therefore, imperils the credibility of the rest of their evidence if readers believe the investigators counter-arguments about the registration of the beacon.

- *Procatalepsis*
  This technique enables an argument to develop by raising and then answering a possible objection. This is intended to avoid a situation in which the reader's attention is distracted from any subsequent argument by the doubts that might have arisen during their reading of the preceding prose. The case study report provides a relatively complex example of this technique: "it is difficult to see how a pilot boat could be completely immune to capsizing or plunging, but pilot boat design criteria must meet the needs of the industry and pilotage authorities" [784]. This illustrates the use of procatalepsis because it addresses the implicit objection that it is impossible to design a pilot boat that is completely immune to capsizing or plunging. The answer to this possible objection is that 'design criteria must meet the needs of the industry and pilotage authorities'. As with previous techniques, this is not without its dangers. The purpose behind the use of procatalepsis is to enable investigators to continue with the main thrust of their argument. There is a danger that such brief comments may do little to address the underlying doubts of the reader. For instance, the Navimar V does not consider the form that such criteria might take not does it address the problems of establishing consensus about the needs of industry and pilotage authorities.

This list presents a preliminary analysis of the rhetorical devices used in incident reports. It builds on initial work by Snowdon [749]. He has argued that it is possible to apply this style of analysis as a means of hecking whether or not particular linguistic constructs are (ab)used to support bias in incident and accident reports. The objective of his work is to teach investigators to perform a detailed and critical reading-through of their reports prior to publication. The intention is not that they should be forced to learn the complex names and ideas associated with each trope. It is, however, intended that greater attention be paid to the effect that particular devices can have upon the readers of an incident report.

Logic

As mentioned, the previous list only provides a partial account of the many rhetorical devices that can be identified in incident and accident reports. Over sixty of these are identified by Harris' Handbook of Rhetorical Devices [307]. Brevity prevents a more sustained analysis. It is, however, worth pausing to consider one additional trope known as *enthymeme*. This is an informally-stated syllogism in which either a premise or the conclusion is omitted. This can be illustrated by the following quotation for the Navimar V case study: "since visibility was good, conduct of both the vessel and the pilot boat was carried out by visual observation during the approach of the two vessels and the transfer of the pilots" [784]. This is an enthymeme because it omits the premise that if visibility is good then such manuevres should be performed using visual observation.

It is also possible to omit the conclusion in an enthymeme if it can be 'generally' understood from the premises. For instance, a meeting was convened between Transport Canada, the Corporation des pilotes du Saint-Laurent central, the Corporation des pilotes du Bas Saint-Laurent and the Laurentian Pilotage Authority to decide upon an initial response to the capsize of the pilot vessel. It was agreed that "Transport Canada would issue a Ship Safety Bulletin if the parties came to a consensus, but such consensus was not reached" [784]. This omits the conclusion that no Ship Safety Bulletin was issued. This illustrates a potential danger that stems from the use on enthymemes. The previous quotation provides no guarantees that such a Bulletin was not issued for other reasons.

In logical terms the 'if' in the preceding extract represents implication not bi-implication. Many readers of this extract may, however, make the inference that the publication was not issued and that this can be entirely explained in terms of the lack of agreement between the parties involved in the investigation.

There are further forms of enthymeme. An initial premise can create a specific context that affects the readers' interpretation of subsequent, more general, premises. Readers may then apply the initial premise to the generalisations in order to infer a number of more particular conclusions. For instance, the following excerpt refers to a range of human factors issues that may have affected the Navimar V incident:

> "The level of care and skill required of a crew manoeuvring a pilot boat are significant factors in this occurrence. Even the most experienced master may suffer a moment's inattention. An emergency manoeuvre to correct the vessel's behaviour may be as harmful as poor vessel design. The human factor is also part of the operating system." [784]

The initial sentence states that operator behaviour contributed to this particular incident. The following sentences provide generalised premises that do not refer directly to the circumstances surrounding the capsize of the Navimar V. The implicit conclusions that readers might identify from this enthymeme is that, in this particular incident, key personnel suffered from a moment's inattention, an emergency maneuver may have taken place and that human factors issues may have impaired the operation of the Navimar V.

There are considerable dangers in the use of enthymemes within incident reports. As we have seen, they often rely upon readers inferring conclusions that are implicit within the premises that appear in the published account. Unfortunately, there are few guarantees that every reader will correctly form the implied syllogism. In particular, the distinction between implication and bi-implication can lead to numerous problems with a negated premise. Statements of the form *not A* and *if A then B* does not enable us to conclude *notB*. However, we can conclude *not B* if we have a premise of the form *A if and only if B*. These problems may appear to be of esoteric significance. They have, however, resulted in numerous objections to the accounts that are presented in incident reports [427]. One solution is to use formal logic as a means of verifying that incident reports present all of the information that readers require in order to form the syllogisms that are used by investigators [412]. It is important to emphasise that there are some important differences between this use of logic and that proposed by Ladkin and Loer [470], reviewed in Chapter 11. WBA uses causal logics to support the causal analysis that must be conducted before a report is drafted. In contrast, the techniques that I have developed are aimed more at improving the presentation and argument in incident reports. Philosophically these differences are important because WBA embodies an objective view of causation from Lewis' approach to counterfactual reasoning [491]. This creates some technical problems when there may be rival explanations for the same observed events; Chapter 11 describes proposals to resolve this by introducing weightings into Lewis' modal structures. In contrast, the use of logic to verify the content of incident reports can avoid making any strong assumptions about what actually was the cause of an incident. This can be left up to the skill and expertise of the investigators. The use of logic in this context is simply intended to ensure that the account of an incident avoids the problems associated with the inappropriate use of enthymemes.

The practical application of logic to support the verification of incident reports is very close to that described in Chapter 9. In this previous chapter, mathematically-based notations were used to reconstruct the events leading to an incident. Rather than constructing clauses from the primary evidence that is obtained in the aftermath of an incident, verification proceeds by building a formal model from the phrases in a report document. This can be illustrated by the previous excerpt from the Navimar V case study. It was agreed that "Transport Canada would issue a Ship Safety Bulletin if the parties came to a consensus, but such consensus was not reached" [784]. This an be formalised using the following clauses, note that some of the parties to the agreement have been omited to simplify the exposition:

$$consensus(transport\_canada, laurentian\_pilotage\_authority) \Rightarrow$$

$$issue(transport\_canada, ship\_safety\_bulletin). \tag{13.1}$$

$$not\ consensus(transport\_canada, laurentian\_pilotage\_authority). \tag{13.2}$$

The first clause states that if consensus is reached between Transport Canada and the Laurentian Pilotage Authority then Transport Canada issues a Ship Safety Bulletin. The second clause states that consensus was not reached. Unfortunately, the laws of first order logic do not enable us to make any inferences from these premises about whether or not a bulletin was issued. This can be illustrated by the following inference rule that represents arguably represents the informal inferences that many readers would apply to the previous quotation. If we know that A is true and that if A is true then B is true, we can conclude that B is indeed true:

$$A, A \Rightarrow B \vdash B \tag{13.3}$$

Unfortunately, this rule cannot be applied to clauses (13.1) and (13.2) because these take the form $A \Rightarrow B, notA$. In order to address any potential confusion, we would be forced to explicitly state that no bulletin was issued by Transport Canada:

$$not\ issue(transport\_canada, ship\_safety\_bulletin). \tag{13.4}$$

Alternatively, we could re-write the prose used in the incident report: Transport Canada would *only* issue a Ship Safety Bulletin if the parties came to a consensus, but such consensus was not reached. The introduction of the modifier 'only' rules out other circumstances that might have led to the publication of the bulletin and which are not mentioned in that particular passage. This would result in the following formalisation which includes the $\Leftrightarrow$ operator (read as 'if and only if'):

$$\begin{aligned} &consensus(transport\_canada, laurentian\_pilotage\_authority) \Leftrightarrow \\ &\quad issue(transport\_canada, ship\_safety\_bulletin). \end{aligned} \tag{13.5}$$

$$not\ consensus(transport\_canada, laurentian\_pilotage\_authority). \tag{13.6}$$

The use of the $\Leftrightarrow$ operator provides a number of additional inference rules that can be used to verify the informal reasoning process that has been described in previous paragraphs. One of these rules can be formalised as follows:

$$not\ A, A \Leftrightarrow B \vdash not\ B \tag{13.7}$$

The proof proceeds by applying (13.7) to clauses (13.5) and (13.6) to derive:

$$not\ issue(transport\_canada, ship\_safety\_bulletin). \tag{13.8}$$

The use of formal logics offers a number of additional benefits to the verification of incident reports. In particular, it can be used to strip out repetition when it is used as a rhetorical device. For example, the Navimar V case study includes the following phrases:

> "In the compulsory pilotage areas on the St. Lawrence River, most pilots use the accommodation ladder for access to vessels." (Section 1.12.2, [784])
>
> "The custom on the St. Lawrence River for a number of decades has been for pilots to use the accommodation ladder rather than the pilot ladder, unless exceptional conditions require a departure from that practice." (Section 2.2, [784])
>
> "Most St. Lawrence River pilots use the accommodation ladder to board vessels." (Section 3.1, [784])

Such repetition can have an adverse effect on the reader of an incident report. The recurrence of similar sentences reiterates particular observations. This indirectly lends additional weight to arguments even though each restatement of the information is based on the same evidence. There may be insidious effects when, as in the previous examples, no evidence is cited to support particular assertions about existing practices on the St. Lawrence. The previous citations might be represented by the following clauses.

$$\begin{aligned}&\mathit{size\_of}(\mathit{pilot}(P1), \mathit{perform}(P1, \mathit{access\_accommodation\_ladder}), N1) \wedge \\ &\mathit{size\_of}(\mathit{pilot}(P2), \mathit{perform}(P2, \mathit{access\_pilot\_ladder}), N2) \wedge \mathit{most}(N1, N2).\end{aligned} \tag{13.9}$$

As mentioned, such formalisations help to strip out the rhetorical effects of repetition. Logical conjunction is idempotent. In other words, $A \wedge A \wedge A \wedge A$ is logically equivalent to A even though the rhetorical effect may be quite different. It is important to note, however, that we have had to rely upon a second order notation in order to formalise the notion of 'most' in (13.9). This illustrates a limitation of our application of logic. A range of relatively complex mathematical concepts may be required in order to formalise the prose within an incident report. There are further limitations. For example, we have not provided the semantics for *most*. We might have resorted to the use of $>$ but this would not have captured the true meaning of the investigators' remarks. Such a formalisation would evaluate to true if just one more pilot used the accommodation ladder rather than the pilot ladder. We might, therefore, specify that $\mathit{most}(N1, N2)$ is true if $N1$ is twice as big as $N2$, three times as big as $N2$, four times as big as $N2$... The key point here is that the process of formalisation forces us to be precise about the meaning of the prose that is used within an incident report. This offers important safeguards during the verification of a particular document. For example, if we consider precise numerical values to support the definition of *most* we might then require that investigators providing statistical evidence to demonstrate that three, four or five times as many Pilots use the accommodation ladders as use the Pilot ladders.

The previous example has illustrated not only how logic can be used to combat the rhetorical effects of repetition, it has also illustrated the level of precision that this approach promotes during the verification of an incident report. There are further benefits, especially when investigators consider variants of the enthymeme tropes mentioned above. An enthymeme involves the omission of a premise or conclusion from an argument. It is relatively rare to find incident reports that deliberately omit major facts from their account [426]. More frequently, evidence can be cited many pages away from the arguments that it supports This creates problems because readers can easily overlook this confirmatory evidence and hence may not draw the conclusions that might otherwise have been derived about the course and causes of an incident. Similar problems can arise from the use of the structuring mechanisms that have been described in previous sections of this chapter. In particular, by stating the conclusions at the end of an incident report it is relatively easy for readers to forget or overlook the caveats and provisos that may have been used to circumscribe those findings in the previous sections of a report. For example, the Navimar report includes the following conclusion: "9. the pilot aboard the bulk carrier and the master of the pilot boat did not come to an agreement by radio communication on the time and position for the transfer" [784]. An initial reading of this finding might suggest that radio communication ought to have been made and that this might have helped to avoid the incident. Such an interpretation ignores some of the contextual factors that convinced both of these experienced mariners that such a course of action was unnecessary. For instance, the report make the following observations ten or more pages before the conclusions cited above:

> "Since neither the master of the pilot boat nor the pilot on board the bulk carrier was expecting any problems with the transfer manoeuvre, they did not see any point in making contact by radiotelephone to determine when the pilot boat should come alongside and transfer the pilots, nor were they required to do so by regulations." (Section 1.9.2, [784])

In order to correctly interpret finding 9, cited above, readers must remember that the pilot and master were not required to make radio contact and that both considered such contact to be unnecessary given the prevailing conditions at the time of the transfer. This is a non-trivial task. As we have seen, most incident reports contain a mass of contextual detail. For example, the Navimar V report provides details of previous pilot transfers that did not play any direct role in this particular incident.

It can be difficult for readers to identify those details that will be significant to their understanding of the conclusions in an incident report and those that simply add circumstantial information. Logic can be used to strip out this contextual information. we have pioneered a style of analysis that is similar to the WBA, described in Chapter 9. Rather than starting with a temporal sequence of events leading to an incident, we start with the conclusions in an incident report. For example, the conclusion about the lack of communication can be represented by the following clause:

$$notmessage(pilot\_bulk\_carrier, master\_pilot\_boat, transfer\_details). \tag{13.10}$$

The verification process then proceeds by a careful reading of the incident report to identify any previous information that relates to this conclusion. The previous citation might be formalised as part of this analysis in the following manner:

$$\begin{aligned}&weather(visibility\_good) \wedge \\ &notrequired\_message(pilot\_bulk\_carrier, master\_pilot\_boat, transfer\_details) \Rightarrow \\ &notmessage(pilot\_bulk\_carrier, master\_pilot\_boat, transfer\_details).\end{aligned} \tag{13.11}$$

Ideally, we would like to apply a formal proof rule, such as (13.3), to show that the conclusion was supported by available evidence. In order to do this we must first demonstrate that the visibility was good and that there was no requirement for the master and the pilot to communicate the details of the transfer. The report contains detailed meteorological information: "since visibility was good, conduct of both the vessel and the pilot boat was carried out by visual observation during the approach of the two vessels and the transfer of the pilots" (Section 1.8.2, [784]). Much less information is presented about the regulatory requirements, which might otherwise have required that the communication take place. This example illustrates the way in which logic can be used to focus on particular aspects of an incident report. The evidence that supports particular conclusions can be described in a precise and concise manner. This directs further analysis of an incident report. Investigators must identify those passages that provide the evidence to support these conclusions. Such extracts can, in turn, be translated into a logic notation to complete the formal proof in a manner similar to that illustrated in previous paragraphs.

A number of caveats can be raised about our use of logic to verify particular properties of an incident report. In particular, our formalisations have relied upon relatively simple variants of first order logic. These lack the sophistication of the more complex, causal techniques that have been derived from Lewis' work on counterfactual arguments [491]. This issue is discussed at greater length in Chapter 9. As we shall see, however, there are more fundamental objections against the use of any logic formalism to analyse the arguments that are put forward within an incident or accident report [775]. For now, however, it is sufficient to briefly summarise a number of additional benefits that can be derived from this technique. In particular, it is possible to demonstrate inconsistencies between two or more accounts of the same incident [427]. For instance, if one report omits a particular piece of evidence then we must consider not only the impact of that omission itself but we must also account for the loss of any inferences that may depend upon that evidence. In the previous example, if a rival report failed to provide any information about the prevailing weather conditions then readers might doubt clause (13.11) as an 'explanation' for the lack of communication. It is also possible to extend the formal model of an incident report to analyse any proposed recommendations. For instance, in previous work we have used formal proof techniques to show that proposed interventions following a rail collision need not prevent the recurrence of a similar incident in the future [427].

Toulmin

Toulmin's 'The Uses of Argument' [775] can be seen as a measured attack against the use of formal logic as a primary means of understanding rational argument. This work, therefore, has considerable importance for any attempt to use logic as a means of verifying the correctness and consistency of arguments within incident reports. One aspect of Toulmin's attack was that many arguments do not follow the formal rules and conventions that are, typically, used to construct logics. An example of this is the use of warrant, or an appeal to the soundness criteria that are applied within a particular field of argumentation. These criteria differ between fields or domain. The style of warrant that might be acceptable within a court of law might, therefore not be acceptable in a clinical environment. Similarly, the clinical arguments that support a particular diagnosis are unlikely to exploit the same soundness criteria that might establish validity within the domain of literary criticism. The abstractions of a formal logic are unlikely to capture the argumentation conventions that characterise particular domains. Toulmin urges us to question the notions of objective or 'universal' truth. The truth of a statement relies upon the acceptance of a set of rules or procedures that are accepted within a domain of discourse.

This initial analysis of Toulmin's work can be applied to the verification of incident reports. For instance, it is possible to identify certain norms and conventions within particular reporting systems. These norms and conventions help to define what is and what is not an acceptable argument about causes of an incident. For example, eye-witness testimony typically provides insufficient grounds for a causal analysis unless it is supported by physical evidence. Conversely, the data from a logging system is unlikely to provide a sufficient basis for any argument about the systemic causes of an adverse occurrence or near-miss incident. One approach to the verification of incident reports would be to enumerate a list of these conventions that are often implicit assumptions within an investigation team. Any report could then be checked against this list to confirm that it met the appropriate argumentation conventions. One of the strong features of this field-dependent aspect of Toulmin's work is that it reflects the differing practices that are apparent between different reporting systems. What might be acceptable as a valid argument about causation for a local investigation into a low-risk incident might not be acceptable within a full-scale investigation of a high-risk incident. Similarly, the standards that might be applied to the argument in an incident report within the nuclear industry might be quite different from those that would be acceptable within a catering business. This domain-dependent approach contrasts strongly with logic-based techniques that rely upon formal notions of correctness. The underlying proof procedures remain the same irrespective of the domain under investigation. This has important practical consequences. The high costs that can be associated with the application of formal modelling techniques can prevent it from being used to verify the correctness of low-risk incidents. The complexity of higher-risk failures can also force investigators to construct abstract models of the information that is contained in accident reports. These models often fail to capture important details of an incident or accident. In either case, problems are apparent because of the inflexibility of logic-based techniques. They cannot simply be tailored to reflect the different forms of argumentation that are exploited within different contexts. Some people would argue that this is a strong benefit of a formal approach; it avoids the imprecision and inconsistencies that can arise from domain dependent argumentation procedures.

The conflict between logic-based models and Toulmin's ideas of domain-dependent discourse has had a profound impact upon the theory of argumentation. Fortunately, many of the consequences of Toulmin's ideas do not apply within the domain of incident reporting. In particular, large-scale reporting systems often encourage their investigators to adopt a model of argument in their reports that closely mirrors aspects of formal proof. Evidence is presented in a reconstruction section, arguments are then developed within an analysis section, conclusions are then presented on the basis of the arguments. It can, therefore, be argued that the domain dependent procedures of incident and accident investigation are similar to our previous application of proof procedures such as Modus Ponens. This, in turn, explains why so many people have proposed logic based techniques as appropriate means of verifying the products of incident investigations [469, 412].

Toulmin acknowledges that people rely upon both domain-dependent and domain-independent procedures to establish the validity of an argument. In contrast to the field-dependent issues mentioned in the previous paragraphs, , most of Toulmin's work focuses on domain-independent procedures. The simplest of these procedures consists of a claim that is supported by some data. A claim is an assertion, for example about a cause of an incident. There are three types of claim:

1. Claims of fact. A claim of fact is supported by citing data, such as the results of simulator studies or of data recorders. This data must be sufficient, accurate, recent, appropriate.

2. Claims of value. These claims represent moral or aesthetic judgements which are not factual and cannot be directly supported by data alone. They can be supported by citing unbiased

and qualified authorities. A claim of value can also be supported by arguing that it produces good results or that negative results may be obtained if it is ignored.

3. Claims of policy. A policy claim is supported by showing that a procedure of regulation is both feasible and positive. Such arguments tend tpo rely upon a combination of fact and of value.

Most incident reports rely upon claims of fact. Arguments about the causes of an incident must be grounded in the evidence that can be obtained by primary and secondary investigations. We are also often concerned with claims of policy. For instance, there may be little a priori evidence that a particular recommendation will avoid or mitigating future incidents. The best that can be done is to follow the advice of relevant experts based on data from similar systems. It is more rare for an incident report to be concerned with value claims, except in circumstances where moral decisions must be made, for instance, over the amount of money that might be invested to avoid future fatalities. as mentioned, incident reports are primarily concerned with claims of fact. The rest of this section, therefore, focuses on the arguments that can be used to support this form of argument.

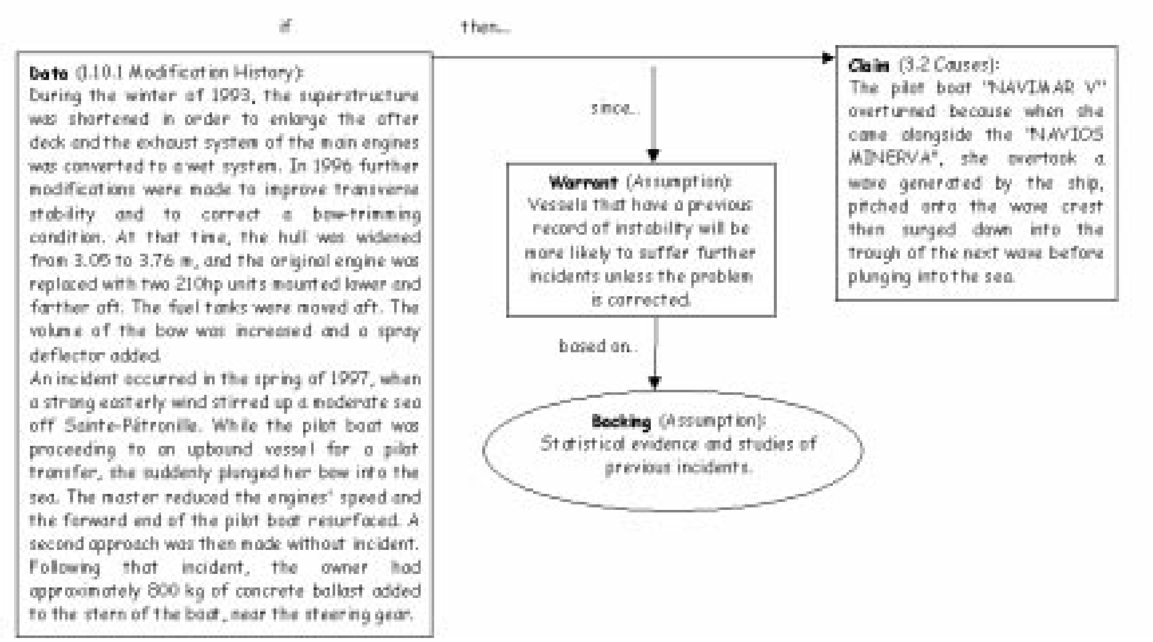

Data (1.10.1 Modification History):
During the winter of 1993, the superstructure was shortened in order to enlarge the after deck and the exhaust system of the main engines was converted to a wet system. In 1996 further modifications were made to improve transverse stability and to correct a bow-trimming condition. At that time, the hull was widened from 3.05 to 3.76 m, and the original engine was replaced with two 210hp units mounted lower and farther aft. The fuel tanks were moved aft. The volume of the bow was increased and a spray deflector added.
An incident occurred in the spring of 1997, when a strong easterly wind stirred up a moderate sea off Sainte-Pétronille. While the pilot boat was proceeding to an upbound vessel for a pilot transfer, she suddenly plunged her bow into the sea. The master reduced the engines' speed and the forward end of the pilot boat resurfaced. A second approach was then made without incident. Following that incident, the owner had approximately 800 kg of concrete ballast added to the stern of the boat, near the steering gear.

Warrant (Assumption):
Vessels that have a previous record of instability will be more likely to suffer further incidents unless the problem is corrected.

Claim (3.2 Causes):
The pilot boat "NAVIMAR V" overturned because when she came alongside the "NAVIOS MINERVA", she overtook a wave generated by the ship, pitched onto the wave crest then surged down into the trough of the next wave before plunging into the sea.

Backing (Assumption):
Statistical evidence and studies of previous incidents.

Figure 13.2: Data and Claims in the Navimar Case Study.

As mentioned, data can be used to support particular claims of fact. Data represents a body of evidence that can be used to determine whether or not a claim is valid. This concept is not straightforward because there can be further arguments about the validity of an item of evidence. It is for this reason that data, or grounds, refers to the part of an argument that is not in dispute. Toulmin avoids some of the practical issues that can arise when the different parties to an incident investigation cannot agree about what is and what is not acceptable evidence:

> "Of course we may not get the challenger to even to agree about the correctness of these facts, and in that case we have to clear his objection out of the way by a preliminary argument: only when this prior issue or 'lemma', as geometers would call it, has been dealt with, are we in a position to return to the original argument. But this complication we need only mention: supposing the lemma to have been disposed of, our question is how to set the original argument out most fully and explicitly". [775]

Chapters 6 and 7 describe techniques that are intended to encourage agreement over the reliability and accuracy of particular items of evidence. In contrast, Figure 13.2 illustrates how Toulmin approach can be applied to the Navimar case study. As can be seen, the Transportation Safety Board of Canada report argues that the incident was caused when the pilot boat pitched onto a wave crest and surged into the trough of the next wave. This claim of fact is supported in the incident report by evidence of a similar, previous incident in 1997. Figure 13.2 also introduces two further components of the Toulmin model of argumentation. A *warrant* describes the assumptions that help to connect a claim with the grounds or data that supports it. This illustrates a superficial relationship between Toulmin and the syllogisms that have been introduced in previous sections [305]. The grounds and warrant can be though of as premises, the claim represent the conclusion to be drawn. Figure 13.2 also illustrates the notion of *backing*. This helps to establish a warrant; backing can also be a claim of fact or value. In this example, the relationship between the previous incident and the conclusion of the report is never made explicit in the incident report. In consequence, readers have to infer the reason why this data might support the overall conclusions. Vessels that have suffered previous incidents are more likely to suffer future recurrences of similar incidents. This warrant might be supported by statistical studies to indicate that vessels which are involved in one incident and then more likely to ve involved in another similar incident in the future.

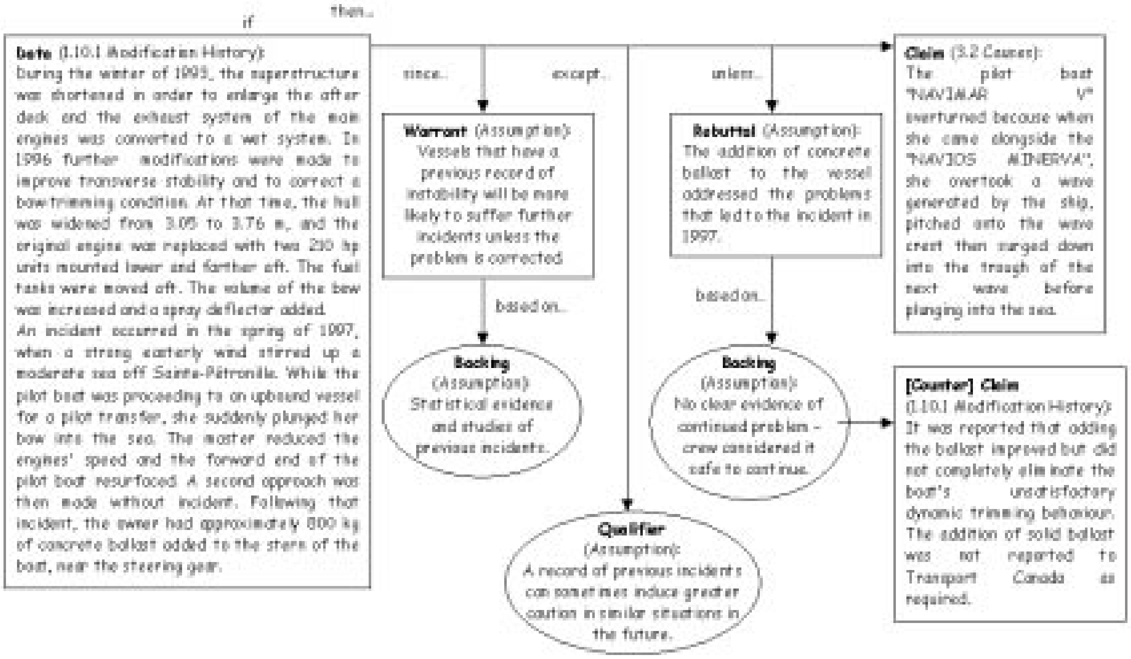
if
then...
Data (1.10.1 Modification History): During the winter of 1995, the superstructure was shortened in order to enlarge the after deck and the exhaust system of the main engines was converted to a wet system. In 1996 further modifications were made to improve transverse stability and to correct a bow-trimming condition. At that time, the hull was widened from 3.05 to 3.76 m, and the original engine was replaced with two 210 hp units mounted lower and farther aft. The fuel tanks were moved aft. The volume of the bow was increased and a spray deflector added.
An incident occurred in the spring of 1997, when a strong easterly wind stirred up a moderate sea off Sainte-Pétronille. While the pilot boat was proceeding to an upbound vessel for a pilot transfer, she suddenly plunged her bow into the sea. The master reduced the engines' speed and the forward end of the pilot boat resurfaced. A second approach was then made without incident. Following that incident, the owner had approximately 800 kg of concrete ballast added to the stern of the boat, near the steering gear.
since...
except...
unless...
Warrant (Assumption): Vessels that have a previous record of instability will be more likely to suffer further incidents unless the problem is corrected.
Rebuttal (Assumption): The addition of concrete ballast to the vessel addressed the problems that led to the incident in 1997.
based on...
based on...
Backing (Assumption): Statistical evidence and studies of previous incidents.
Backing (Assumption): No clear evidence of continued problem - crew considered it safe to continue.
Qualifier (Assumption): A record of previous incidents can sometimes induce greater caution in similar situations in the future.
Claim (3.2 Causes): The pilot boat "NAVIMAR V" overturned because when she came alongside the "NAVIOS MINERVA", she overtook a wave generated by the ship, pitched onto the wave crest then surged down into the trough of the next wave before plunging into the sea.
[Counter] Claim (1.10.1 Modification History): It was reported that adding the ballast improved but did not completely eliminate the boat's unsatisfactory dynamic trimming behaviour. The addition of solid ballast was not reported to Transport Canada as required.

Figure 13.3: Qualification and Rebuttal in the Navimar Case Study.

An argument is valid if the warrant and any associated data provides adequate support for the claim. It is possible, however, that can investigators' colleagues might raise objections to a particular argument. Alternatively, the writer of an incident report may themselves have doubts about the scope and applicability of their analysis. Finally, the readers of an incident report might question the argument that is embodied within such a document. Toulmin's domain independent model can be used to capture these alternate positions that challenge or modify an initial position. Figure 13.3 uses the Navimar case study to illustrate such an extension. As can be seen, a qualification node had been introduced to record the observation that a previous incident can induce greater caution amongst the crew of some vessels. This qualifier effects the previous warrant the represents the implicit argument that vessels involved in previous incidents will be more likely to be involved in future incidents. The claim can also be qualified in a similar fashion, if it has been challenged and its truth is in doubt. This use of a qualifier does not deny that there may be a relationship between the vessels involved in an incident and previous involvement in similar occurrences. Instead, it argues that this effect may not apply to all vessels.

Figure 13.3 also illustrates the use of a rebuttal to challenge the argument that was first sketched in Figure 13.2. The suggestion that vessels are more likely to be involved in an adverse occurrence if they have been involved in previous incidents is contradicted in this case by arguing that previous modifications were effective in addressing the causes of the problem. This rebuttal is supported by the lack of evidence to indicate that there was a continuing problem. The crew also continued to risk their lives by operating the vessel in what can be a difficult and hazardous environment. Such counter-arguments can be challenged. For instance by arguing that economic pressures often force individuals to operate equipment that they know to have safety problems. Similarly, the lack of evidence about further incidents between spring 1997 and the incident need not provide direct assurance that the problem would not recur. The vessel might not have faced similar operating conditions. Such arguments against a rebuttal could be incorporated into the structures of Figures 13.2 and 13.3. The resulting graphs can quickly become intractable. A number of researchers have, therefore, developed tools that can be used to support the use of Toulmin's techniques to map out argument structures [503, 748]. Much of this work has been inspired not by Toulmin's initial interest in studying the structure of argument but by a more prosaic interest in improving the support that is provided for particular decisions. This practical application of Toulmin's model can be illustrated by Locker's recent work on improving 'Business and administrative communication' [498]. Locker suggests that business writers decide how much of Toulmin's model they should use by analysing the reader and the situation. Writers should make both their claim and the data explicit unless they are sure that the reader will act without questioning a decision. The warrant should be included in most cases and the backing should be made explicit explicit. It is also important for effective communication that any rebuttals should be addressed by counter-claims, as suggested above. Authors should also be careful to explicitly address any limiting or qualifying claims.

Locker's normative application of the Toulmin model suggests how this approach might be used to support the presentation of incident and accident reports. Given the importance of effective communication in this context, we might require that this approach be used to make explicit the association between data and the claims that are made in a report. If data is not presented in the document then the claim is unsupported. This application of the technique is similar to the manner in which logic might be used to identify enthymemes in other forms of syllogism. It can also be argued that investigators should explicitly document the warrant that links data and evidence within an incident report. This is important because, as we have seen, incident reports are often read by many diverse communities including operators, managers, regulators, engineers etc. It is, therefore, difficult to make strong assumptions about the background knowledge that is required in order to infer the relationships that exist between data and particular clausal claims. This is illustrated by the Navimar case study that has been introduced in the previous paragraphs. The Transportation Safety Board report never makes explicit the relationship between evidence of a previous incident and the overall conclusion of the enquiry. We have had to infer an appropriate warrant in Figure 13.2. It is entirely possible that we have made a mistake. Investigators may have had entirely different reasons for introducing the events in 1997. Unfortunately, we have know way of telling whether this argument is correct or not from the report into the subsequent capsize.

This requirement to make explicit data, claims and warrant goes beyond Locker's requirements for effective business communication. These differences should not be surprising, given that Toulmin acknowledges the importance of domain independent and domain dependent requirements for effective argument. The 'standards of proof' are potentially higher in the case of incident investigation than they might be amongst more general business applications. Other aspects of good practice will be common across these different domains. For example, Locker argues that effective communication relies upon writers explicitly addressing any proposed rebuttals of an initial argument. This is equally important within the field of incident reporting. For instance, the Navimar report comments that 'it was reported that adding ballast improved but did not completely eliminate the boat's unsatisfactory trimming behaviour' [784]. This partly addresses the rebuttal in Figure 13.3. It does not, however, explicitly provide backing for such a counter-argument beyond the rather vague reference to previous reports.

Figure 13.4 provides a slightly more complex example of the application of Toulmin's model to the Navimar case study. In this instance, data about the layout of the pilot vessel is used to support a claim that the use of the accommodation adder rather than the pilot ladder contributed to the incident. This argument is supported by the warrant that the use of the accommodation ladder complicated the task of keeping the pilot boat's transfer deck in position because their boat cannot rest parallel to the vessel's side during the transfer. The difficulty of performing such maneuvers is recognised by international regulations requiring that pilot ladders be used for pilotage transfers. The argument is also supported by the warrant that the master had to divide his attention between completing this relatively complex manuevre with the vessel in front of him and the task of looking aft to ensure that the after deck lined up under the accommodation ladder. This warrant is not backed by any particular citations in the incident report. There are no direct observations or accounts of the difficulty of this task. In contrast, the investigators acknowledge that "there is every indication that the crew were well rested and highly experienced" [784]. The omission of any backing provides a further example of an enthymeme trope. If we apply Toulmin's model in the normative manner proposed above then it can be argued that more information ought to be introduced into the report to support this warrant. For instance, an analysis of the ergonomics of the bridge design might provide sufficient detail for operators to determine whether similar problems might affect not simply transfer tasks but other pilot operations as well.

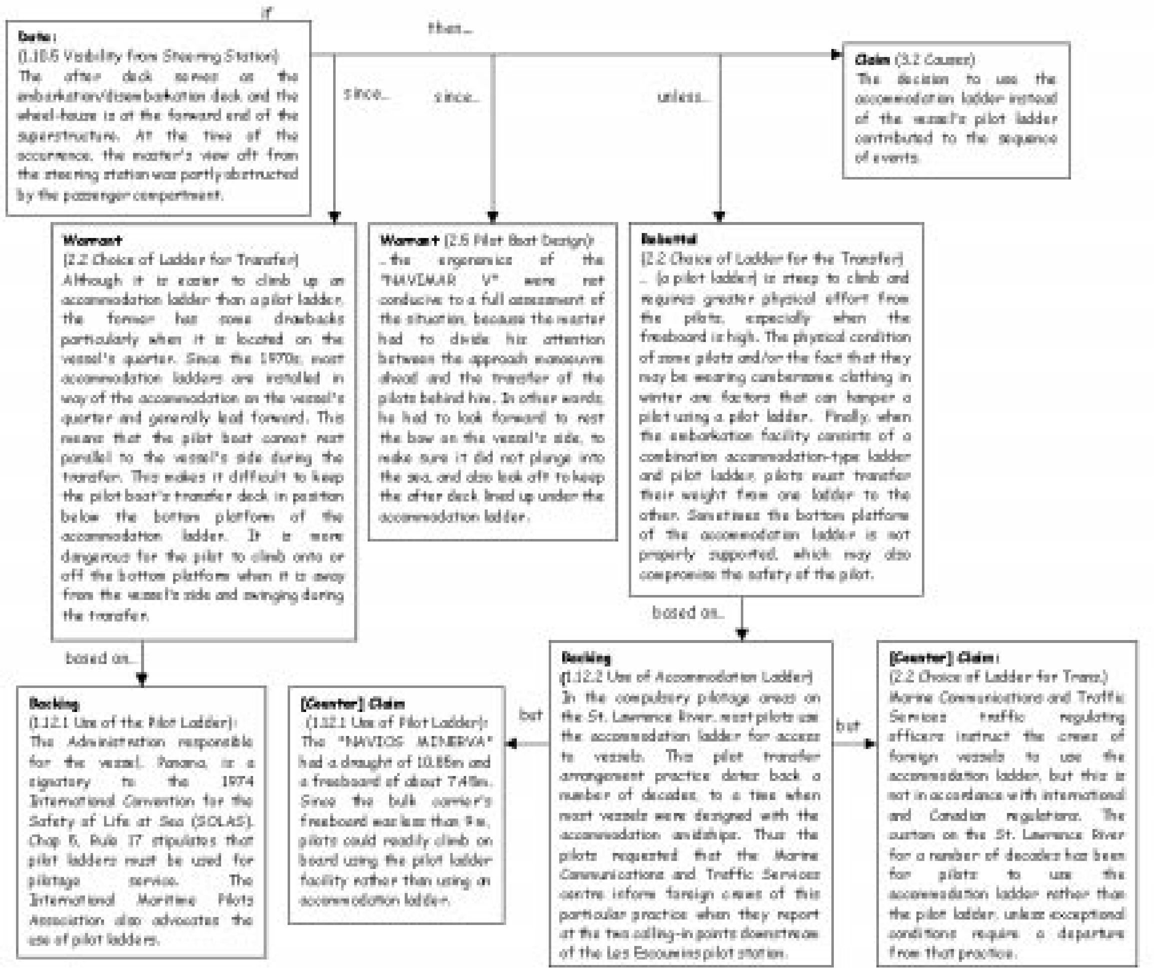

Data:
(1.10.5 Visibility from Steering Station)
The after deck serves as the embarkation/disembarkation deck and the wheel-house is at the forward end of the superstructure. At the time of the occurrence, the master's view aft from the steering station was partly obstructed by the passenger compartment.

Claim (3.2 Causes)
The decision to use the accommodation ladder instead of the vessel's pilot ladder contributed to the sequence of events.

Warrant
(2.2 Choice of Ladder for Transfer)
Although it is easier to climb up an accommodation ladder than a pilot ladder, the former has some drawbacks particularly when it is located on the vessel's quarter. Since the 1970s, most accommodation ladders are installed in way of the accommodation on the vessel's quarter and generally lead forward. This means that the pilot boat cannot rest parallel to the vessel's side during the transfer. This makes it difficult to keep the pilot boat's transfer deck in position below the bottom platform of the accommodation ladder. It is more dangerous for the pilot to climb onto or off the bottom platform when it is away from the vessel's side and swinging during the transfer.

Warrant (2.5 Pilot Boat Design):
...the ergonomics of the "NAVIMAR V" were not conducive to a full assessment of the situation, because the master had to divide his attention between the approach manoeuvre ahead and the transfer of the pilots behind him. In other words, he had to look forward to rest the bow on the vessel's side, to make sure it did not plunge into the sea, and also look aft to keep the after deck lined up under the accommodation ladder.

Rebuttal
(2.2 Choice of Ladder for the Transfer)
... (a pilot ladder) is steep to climb and requires greater physical effort from the pilots, especially when the freeboard is high. The physical condition of some pilots and/or the fact that they may be wearing cumbersome clothing in winter are factors that can hamper a pilot using a pilot ladder. Finally, when the embarkation facility consists of a combination accommodation-type ladder and pilot ladder, pilots must transfer their weight from one ladder to the other. Sometimes the bottom platform of the accommodation ladder is not properly supported, which may also compromise the safety of the pilot.

Backing
(1.12.1 Use of the Pilot Ladder):
The Administration responsible for the vessel, Panama, is a signatory to the 1974 International Convention for the Safety of Life at Sea (SOLAS). Chap 5, Rule 17 stipulates that pilot ladders must be used for pilotage service. The International Maritime Pilots Association also advocates the use of pilot ladders.

[Counter] Claim
(1.12.1 Use of Pilot Ladder):
The "NAVIOS MINERVA" had a draught of 10.85m and a freeboard of about 7.45m. Since the bulk carrier's freeboard was less than 9m, pilots could readily climb on board using the pilot ladder facility rather than using an accommodation ladder.

Backing
(1.12.2 Use of Accommodation Ladder)
In the compulsory pilotage areas on the St. Lawrence River, most pilots use the accommodation ladder for access to vessels. This pilot transfer arrangement practice dates back a number of decades, to a time when most vessels were designed with the accommodation amidships. Thus the pilots requested that the Marine Communications and Traffic Services centre inform foreign crews of this particular practice when they report at the two calling-in points downstream of the Les Escoumins pilot station.

[Counter] Claim:
(2.2 Choice of Ladder for Trans.)
Marine Communications and Traffic Services traffic regulating officers instruct the crews of foreign vessels to use the accommodation ladder, but this is not in accordance with international and Canadian regulations. The custom on the St. Lawrence River for a number of decades has been for pilots to use the accommodation ladder rather than the pilot ladder, unless exceptional conditions require a departure from that practice.

Figure 13.4: More Complex Applications of Toulmin's Model.

Figure 13.4 also provides a further illustration of a rebuttal that readers might form from the information that is presented in the Navimar report. There are a number of drawbacks that affect the use of pilot ladders. Some if these potential problems relate to significant safety concerns, especially if hybrid pilot and accommodation ladders are joined together. This rebuttal is supported by the observation that it was common practice to use accommodation ladders in the pilotage areas of the St. Lawrence River. This had reached such an extent that the Marine Communications and Traffic Service centre explicitly informed foreign crews of this practice. Previous paragraphs have argued that it is important for investigators to address such rebuttals if readers are to have confidence in the findings of an incident report. This caveat is, therefore, addressed by two counter claims. Firstly, the free board of the Navios Minerva was less than nine meters. This obviated the need for a combination ladder of the type mentioned above. The rebuttal is also addressed by the counter claim that the Marine Communications and Traffic Service instruction was neither in accordance with international nor Canadian regulations. The safety concerns mentioned in the initial rebuttal cannot avoid the conclusion that common practice was in violation of the recommended rules and regulations that had been issued to the crews.

A number of further comments can be made about the use of Toulmin's model in Figure 13.4. It is possible to use the resulting graphs to trace the location of information to various sections within the report. Although most of the information that supports the rebuttal appears together in Section 2.2, some of the material can be found in 1.12.2. This is important because this information provides evidence that can be used to contradict the initial rebuttal. Similarly, the backing for part of the warrant, in Section 2.2, supporting the argument in Figure 13.3 is to be found in Section 1.12.1. This forms part of a more general pattern that can be observed through the application of Toulmin's model to incident reports. The warrant that outlines a particular line of support for an argument, typically, appears many pages after the backing that supports it. This separation arises from the policy of separating the arguments and analysis that explain the significance of key events from the reconstructions that first describe the context in which an incident occurs. This approach helps to avoid any confusion between what is known and what is inferred about a near miss incident or adverse occurrence. A further consequence of this is that readers may only learn the significance of particular items of information after they have finished reader the report. In consequence, it is often necessary to read and re-read such documents several times in order to follow the complex argument that may be distributed across hundreds of pages of prose. Snowdon has argued that these problems might be reduced if, instead of using Toulmin to check an argument in an incident report, the readers of a report were provided with diagrams such as those shown in Figures 13.2, 13.3 and 13.4, [748]. These could be printed inside an incident report to provide readers with a 'roadmap' of the various arguments that are being proposed by an investigator. This approach might also increase confidence in any conclusions by explicitly indicating the counter-arguments that might be deployed against particular rebuttals.

Figure 13.4 illustrates the relatively complex diagrams that can emerge from the application of Toulmin's techniques to incident reports. It also illustrates some of the problems that arise in the practical use of this approach. Toulmin's focus was on explaining the domain dependent and domain independent components of argument structures. His purpose was "to raise problems, not to solve them; to draw attention to a field of inquiry, rather than to survey it fully; and to provoke discussion rather than to serve as a systematic treatise" [775]. His model was never intended to be used as a tool to support the development of incident reports. One consequence of this is that is can be difficult to categorise the paragraphs within an incident report. For example, it can be argued that the rebuttal in Figure 13.4 might be reclassified as a form of qualifier. It does not directly contradict the argument that the decision to use the accommodation ladder contributed to the incident. In contrast, it explains why many operators chose not to use pilot ladders. This could be interpreted as a qualifier because it refers to previous instances in which the use of the accommodation ladder had not resulted in an adverse outcome. Further problems complicate this application of the Toulmin model. For instance, we have constructed our analysis at the level of individual paragraphs within the Navimar report. Rather than translate the original prose in a manner that might support the classification of those paragraphs within the Toulmin approach, we have chosen to retain verbatim quotations within our analysis. The drawback to this application of the model is that some of the paragraphs may themselves contain more detailed argument structures. For example, the backing for the rebuttal in Figure 13.4 contains a claim that most pilots use the accommodation ladder in the St Lawrence River pilotage areas. This might, in turn, be supported by additional data. The introduction of this data would then need to be supported by an appropriate warrant and so on.

The previous objections relate to the difficulty of applying Toulmin's model to the complex prose and argument structures that are used in many incident reports. It is possible to reverse these objections by arguing that this very complexity increases the importance of any techniques that might be able to identify potential errors of omission and commission. The principle benefit of this approach is that it provides a graphical representation of various positions within an incident report. These diagrams can then form a focus for subsequent discussion amongst an investigation team prior to publication. The very accessibility of these diagrams helps to ensure that any disagreements about the classification of particular sentences and paragraphs can be checked by an investigator's colleagues. It also forms a strong contrast with the use of more formal logic-based approaches. The accessibility of the Toulmin model, however, comes at the price of far weaker concepts of proof or correctness. The adequacy of an argument can only be assessed in terms of the domain dependent procedures that are accepted within an investigation team. These procedures guide the normative application of Toulmin's model, proposed for business communication by Locker [498] and sketched for incident reporting in previous paragraphs. Problems can arise when those procedures that are acceptable within one domain of argument are questioned or rejected by other groups who employ different standards of 'correctness' .

A number of authors have proposed further extensions to the Toulmin model of argumentation. For instance, the initial proposals provide backing for the propositions that are captures in a warrant. They did not provide similar support for the data that backs a claim. As mentioned above, data is assumed to be accepted. If it is questioned then it must be addressed by a secondary argument. In contrast, Ver Linden has argued that the Toulmin model should be expanded to include 'verifiers' for data [494]. Explicit verifiers involve a further argument, which concludes that the data in the original argument is correct. Thee include citations as well as reference to common knowledge and to personal background. Implicit verifiers stem from the observation that arguers often do not express a clear rationale for accepting data. Instead, people provide a range of cues that are intended to convince the recipient that data is correct. Ver Linden argues that such "sincerity cues probably differ from culture to culture and in the general American culture they include the use of eye contact, tone of voice, and facial expression and other signs of emotion appropriate to the subject, as well as language that emphasises the arguer's sincerity". Implicit verifiers are suggested by the person supporting the claim. In contrast, inferred verifiers are supplied by the receiver without explicit suggestion by the claimant. For instance, the reader of an incident report may believe in an assertion simply because it has been made by a national transportation safety board. It is important to emphasise that this does not imply that the arguers are unaware of the likelihood that recipients will form such inferences. For instance, a national transportation safety board might make an assertion and state it as a fact without citing any source. In such circumstances, they rely on the belief that their reputation will carry a particular weight with the intended audience. This analysis has important implications for the presentation and analysis of incident reports. Data should be supported by explicit verifiers rather than the suggestive expressions of belief provided by implicit verifiers or any reliance on reputation to support the use of inferred verifiers. Ver Linden's analysis not only applies to the backing for data, it can also be applied to the backing that supports warrants and rebuttals. For instance, the backing for the rebuttal in Figure 13.4 clearly relies upon an inferred verifier because no evidence is supplied to support the observation that most pilots use accommodation ladders.

The previous paragraph focused on practical extensions to the Toulmin model. A number of other authors have raised more theoretical objections to this approach. For instance, Freeman introduces the notion of 'gappiness' between a warrant and some data [280]. Warrants can be thought of as inference rules that allow us to move from data to a claim. They are only necessary because the reader senses a gap between the data and the claim that it supports. Freeman's revision has disturbing implications. It can be difficult to predict where readers might sense a 'gap' in an underlying argument. Investigators might, therefore, attempt to exhaustively addresses all of the possible doubts that a skeptical reader might have about an incident report. Such an approach raises further problems. Many supporting arguments would be unnecessary. They exhaustive approach would address gaps that would never occur to many of the readers of an incident report. For instance, it might be necessary to justify the reference to SOLAS in Figure 13.4. The importance and relevance of such agreements would be self-evident to all domain experts. The introduction of additional warrants would support a minority of readers but it would also increases the size and scope of incident reports.

To summarise, Toulmin argues that it is possible to identify domain dependent procedures that help to define convincing arguments within a particular context. Freeman argues that for structural reasons, some of those procedures relate to the individual reader's perception of gaps between claims and data. This implies that normative techniques, such as those proposed by Locker [498], may fail to identify the individual information needs of particular readers. There are currently two practical means that can be used to address Freeman's more theoretical caveats. One technique involved the use of user-testing and experimental analysis to determine whether or not domain dependent procedures are sufficient for a broad cross-section of the intended readership of an incident report. This approach is described in more detail in the next section. It will not address the individual information needs identified by Freeman. It can, however, increase confidence that the argument in an incident report provides sufficient backing to convince a specified proportion of its intended audience.

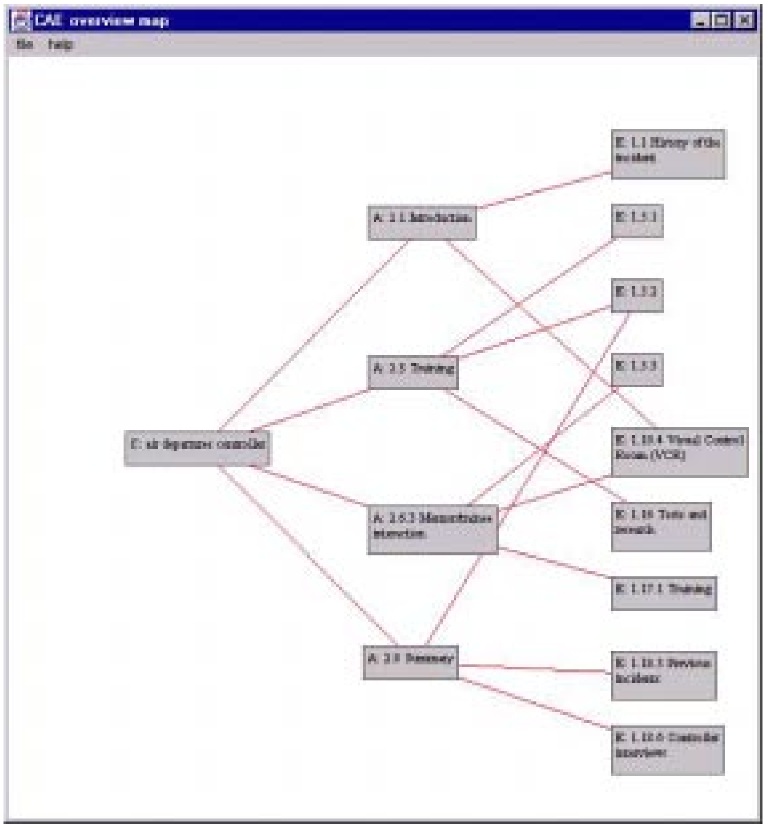
CAE overview map
file
help
C: air departure controller
A: 2.1 Introduction
A: 2.5 Training
A: 2.6.3 Man-machine interaction
A: 2.9 Summary
E: 1.1 History of the incident
E: 1.5.1
E: 1.5.2
E: 1.5.3
E: 1.10.4 Visual Control Room (VCR)
E: 1.16 Tests and research
E: 1.17.1 Training
E: 1.18.3 Previous incidents
E: 1.18.6 Controller interviews

Figure 13.5: Snowdon's Tool for Visualising Argument in Incident Reports (1).

Freeman's caveats about the individual perception of 'gappiness' in argument structures can also be addressed by tool support. These techniques enable readers to view the argument that supports an incident report at a number of different levels of granularity. For example, Snowdon has developed a tool that is based on the Conclusion, Analysis and Evidence structures that were in introduced in Chapter 9. This is a simplification of the full Toulmin model that is specifically intended to support the analysis of incident and accident reports [415]. When the reader of an incident report initially uses the tool, they are presented with a simple overview of the highest level argument. This typically consists of a node that lists the conclusions of the report. By clicking on one of those conclusions they can expand their view of the argumentation in the report to look at the evidence or data that supports a conclusion. By clicking on that data, they can expand their view of the warrant, or analysis, that connects the evidence to the conclusion. This interactive process continues until the user reaches the bottom level in the system. These leaf nodes represent the paragraphs of prose that have been written by the investigators. Figure 13.5 illustrates a partial expansion of the argument structure in an aviation incident report. Figure 13.6 represents the end result of continuing to ask for more information about the investigator's argument. The reader is free to continue to ask for more justification until they reach the sections of prose written by the investigators.

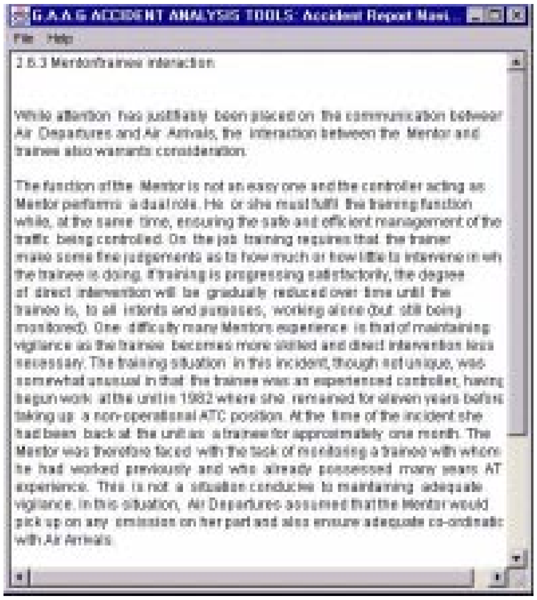

2.6.3 Mentor/trainee interaction

While attention has justifiably been placed on the communication betweer
Air Departures and Air Arrivals, the interaction between the Mentor and
trainee also warrants consideration.

The function of the Mentor is not an easy one and the controller acting as
Mentor performs a dual role. He or she must fulfil the training function
while, at the same time, ensuring the safe and efficient management of the
traffic being controlled. On the job training requires that the trainer
make some fine judgements as to how much or how little to intervene in wh
the trainee is doing. If training is progressing satisfactorily, the degree
of direct intervention will be gradually reduced over time until the
trainee is, to all intents and purposes, working alone (but still being
monitored). One difficulty many Mentors experience is that of maintaining
vigilance as the trainee becomes more skilled and direct intervention less
necessary. The training situation in this incident, though not unique, was
somewhat unusual in that the trainee was an experienced controller, having
begun work at the unit in 1982 where she remained for eleven years before
taking up a non-operational ATC position. At the time of the incident she
had been back at the unit as a trainee for approximately one month. The
Mentor was therefore faced with the task of monitoring a trainee with whom
he had worked previously and who already possessed many years AT
experience. This is not a situation conducive to maintaining adequate
vigilance. In this situation, Air Departures assumed that the Mentor would
pick up on any omission on her part and also ensure adequate co-ordinatio
with Air Arrivals.

Figure 13.6: Snowdon's Tool for Visualising Argument in Incident Reports (2).

Snowdon's tool provides an alternative means both of verifying the argument that is presented in an incident report and of presenting the contents of the report to end users. Investigators can explore the graphical hierarchy to ensure that sufficient backing has been provided for key arguments. Readers can also use the graphical interface to rapidly fill any gaps between data and claims. The key benefits from their perspective is that they need only request additional information about those areas of an argument that they perceive to require additional support. Extensions to the system can also log those areas of a report where users repeatedly ask for additional warrants. Such information can help to guide the presentation of future incident reports.

13.3.2 Validation

There is no direct evidence that any of the techniques described in the previous section will contribute to a sustained improvement in the quality of incident reports within complex, real-world applications. These approaches are the product of research initiatives that have been commissioned by regulatory and other government organisations because of the perceived weaknesses in existing reporting techniques. A number of groups, including my own, are using several of the more recent verification techniques as part of an initial 'field trial'. These studies are intended to yield the evidence that many will require before introducing such 'leading edge' techniques. Until such evidence is available there remains a strong suspicion that the more elaborate approaches, such as the use of formal logics, may contribute little beyond what can be achieved through a careful reading of the prose in an incident report. Two principle objections can be raised to this argument:

1. why do so many people criticise the quality of incident reports if careful reading is sufficient to identify the enthymemes and other rhetorical effects that reflect systematic biases and inappropriate assumptions about the potential readership of these documents?

2. careful reading is, typically, conducted by other investigators. Those investigators often do not reflect the broad range of skills and expertise that characterise the intended audience of an incident report. This is significant because, as Freeman suggests [280], the background of the reader can help to determine the sufficiency of an argument. In other words, members of an investigatory organisation may fail to identify the 'gaps' that will be identified by the eventual readership of an incident report.

The following sections explore the substance of these objections. Firstly, we assess empirical work to determine the nature of existing criticisms against the presentation of incident reports. Subsequent sections then describe a range of techniques that have been used to identify particular weaknesses in individual reports. These validation techniques different from the approaches that were described in the previous section because they focus on end-user testing to assess the utility of these documents. In contrast, verification techniques look more narrowly at whether particular documents satisfy a range of technical properties that can, typically, be established without direct user testing.

User Testing

It is surprising how little research has actually been conducted into the presentation of incident and accident reports. The format described in the previous sections of this chapter has remained largely unaltered for at least one hundred and thirty years [357]. Most of the previous studies commissioned by investigatory and regulatory organisations have focussed on recurrence rates to demonstrate that a reporting scheme has had the desired effect. Relatively little work has been conducted to determine whether any observed improvements might be increased by changing the presentation and dissemination of incident reports. This omission led Snowdon to survey a number of safety managers in an attempt to identify particular attitudes towards incident reporting practices. This study raised considerable practical challenges. In particular, it proved difficult to obtain responses by contacting individuals in their workplace. The sensitive nature of their job can be argued to create a justifiable nervousness in replying to surveys that address their relationship with investigatory and regulatory organisations. He was, therefore, forced to issue questionnaires to a random sample of safety managers at a trade conference. Steps were taken to ensure that the delegates were registered participants at the meeting but assurances were also provided to protect the respondents' anonymity. The average age of respondents was 43 years old. The average experience in a field relevant to safety management was 13 years. These constraints limit the generalisations that can be made from the results of this work. Snowdon argues that it can only be seen as a pilot study but that his observations are strongly suggestive of the existing limitations with incident and accident reports.

Question	Positive responses (N= 27)
Do you use accident reports to inform you of design problems?	19
Do you only read reports that you feel are related to your areas?	17
Do you read the whole of the report?	24
Do you assess the conclusions by checking the evidence or other forms of analysis?	22

Table 13.7: Summary of Results from Snowdon's Survey of Accident Reporting Practices [749]

Table 13.7 presents an overview of the results from part of Snowdon's survey. As can be seen, he focuses on attitudes towards the larger-scale documents that report on accidents and high-consequence incidents. It is interesting to note the relatively high number of respondents who claim to read the entire report and who check the conclusions against either the evidence or the analysis that is presented in these documents. He also documents a number of responses to more open questions about the nature of such reports [748]. Some of these responses provide additional evidence for the broad results that are summarised in Table 13.7. For example, one safety manager was able to illustrate their detailed knowledge of a report that had a direct impact upon their working life. The level of detail in the following response is highly indicative both of the checking that this individual had been motivated to perform and of a careful reading of the entire report:

> "They are poorly structured and are often not tailored to their audiences mixed ability to get an overall picture. Many scenarios are hidden in different parts of the report, e.g. in the Watford junction tain accident in paragraph 143/4 it says a signal sighting committee should have been set up before the accident. After the accident this committee suggested that the removal of some trees would increase the sighting distance. In paragraph 171 it says that if the train had braked earlier, i.e. given the removal of the trees, there would not have been an accident." (Cited in [748])

Some of these responses reveal the continuing perception that many of these documents reflect a blame culture: "investigators usually assume that the victim is the sole cause even when the equipment has glaring design defects that entrap the user". Other responses are less easy to interpret. For instance, one safety manager identified the "tendency on the part of the reporter to write report to support his or her conclusions rather than openly evaluate all of the information/evidence that collected about the event." At one level, this seems to reflect an awareness of the conformation bias that is described in Chapter 11. At another level, it is difficult to know how an investigator could consider every aspect of the evidence that is obtained from a primary and secondary investigation. As we have seen, the drafting of an incident or accident report inevitably involves a filtering or selection process.

Snowdon's survey also provides some confirmation of the analysis that is presented in previous sections of this chapter. Ver Linden's [494] comments on the importance of 'verifiers' are illustrated by the following comment; "they present their account as definitive without acknowledging missing evidence or contradictory opinions". Some of the comments also relate to the concept of 'gappiness' proposed by Freeman [280]. One safety manager wrote that "conclusions draw that do not appear to be 'in line' with facts presented; change in level of strength of assertions - report starts with 'it may have been due to' and ends with 'this was because'." Other feedback relates more strongly to the normative application of Toulmin's ideas proposed by Locker [498]. For example, the following quotation emphasises the importance of explicitly addressing both the qualifiers and rebuttals that can limit the scope of an argument:

> "A good report is one w(h)ere all possibilities are considered and a balanced view is taken. If the report is not able to be conclusive then so be it. All possible causes are listed." (Cited in [748])

As with the previous comment about accounting for every item of evidence, this response illustrates the high expectations that many safety managers have for the presentation of information in accident and incident reports. Chapter 11 describes how causal asymmetries prevent investigators from identifying all of the possible causes for any observed set of effects. The results of previous studies into the theoretical and technical foundations of causation have, therefore, had little impact on the practical actitivies of safety managers. This is surprising given that some of the responses to the survey show a considerable level of knowledge about previous studies of accidents and incidents. These studies have clearly provided a vocabulary with which to voice their criticisms:

> "The report just contains direct causes of accident but information of underlying factors of accident is not available. Accident report should encompass not only direct causes but also proximal causes as well as distal causes. It means that failures mechanism should cover not only operative's failure but also management and organisational failure including design failure". (Cited in [748])

The main impression gained from an analysis of these various comments was that many safety managers are unhappy with the structure and format of the information that is presented in incident reports. Several comments related to the problems of using these documents to inform their daily activities: "...the information I am interested in (human error contribution to the accident and cognitive factors in general) is dispersed over different sections of the report, and that related information might be 'hidden' and I will have to spend a lot of time and effort trying to find it". It can be argued, however, that incident and accident reports are not intended to be used in this way. Their primary purpose is to trigger regulatory action by the bodies that commission individual reports. They are not intended to inform local initiatives by local safety managers. Such a response ignores the high level of interest in these documents that is reported by the safety managers that Snowdon surveyed. They seem to perceive these documents as important components of any safety management system that should be, and are, read as part of their normal working activities. The negative reaction to existing reporting techniques can be summarised by the following response:

> "It is becoming less uncommon to find a report that reflects little effort at gathering evidence. Consequently the analysis and conclusions are shallow. Often when the analysis is shallow, the few facts available are over-emphasised, as if the writer knows the facts are insufficient and attempts to cover by reaching firm conclusions." (Cited in [748])

Unfortunately, there are theoretical and practical barrers that prevent investigators from addressing some of the criticisms that were voiced in Snowdon's survey. As we have seen, it is often impossible to accurately summarise all of the evidence that is collected in the aftermath of an adverse occurrence or near miss incident. Similarly, it is infeasible to consider *every possible* cause of a failure. Freeman [280] provides a possible solution when he emphasises the subjective nature of the problems that many readers experience when they read complex documents. He argued that the notion of a 'convincing' argument depends upon the information that a reader requires in order to bridge the devide between a claim and some evidence. It, therefore, follows that any claims to the sufficiency of an argument make little sense without additional validation to assess whether or not the intended audience can make the necessary connections.

User testing provides one means of assessing whether or not particular groups of individuals are convinced by the argument in an incident report. This approach is described in a large number of introductory textbooks. The following pages, therefore, briefly summarise the main features of the available techniques. The interested reader is directed to [686] and [740] for a more sustained analysis. It is possible to distinguish between two different approaches to validation that can be applied to assess the quality of incident reports. Summative techniques can be used at the end of the production phase when a report is ready to be issued. In contrast, formative evaluation techniques can be used to determine whether particular sections of a report provide the necessary feedback the safety managers and regulators need to complete their tasks.

Formative evaluation helps to guide or form the decisions that must be made during the drafting of an incident report. The importance of this form of incremental user testing depends on the scale of the incident report. It also gives rise to an important paradox. Reports into 'higher-risk' incidents are, typically, longer and more complex than lower risk occurrences or near misses. In consequence, there is a greater need to identify potential problems in the presentation of material about the incident. Any potential 'gaps', omissions or inconsistencies should be identified well before the document is published. In contrast, it is precisely these documents that create the greatest concerns about security and media interest prior to publication. These concerns, mentioned in the opening sections of this chapter, create considerable practical problems when recruiting subjects to provide feedback on the information that is contained in a draft report.

In contrast to formative evaluation, summative evaluation takes place immediately prior to the publication of an incident report. This approach can be problematic unless it is supported by other forms of quality control. It can be costly and time-consuming to make major structural changes to the argument that is presented in an incident report at this late stage in the investigation process. User testing can, occasionally, identify rebuttals that are not addressed either by existing arguments or by available evidence. This is particularly the case when investigators may not have the same degree of domain expertise as the individuals reading the report [248]. Summative evaluation can identify these potential doubts at a time when there are insufficient resources available to commission additional studies of the available evidence.

User-testing methods have been widely applied to computer systems and to the documentation that is intended to support them. As we shall see, some of these techniques can be applied to support the validation of incident reports. The following list identifies some of these approaches and indicates whether they offer the greatest benefits for summative or formative evaluations:

- *Scenario-Based Evaluation.* Scenarios or sample traces of interaction can be used to drive both the drafting and evaluation of an incident report systems. This approach forces investigators to identify key requirements of an incident report. These requirements are summarised as brief descriptions of the sorts of prototypical tasks that different readers might want to perform with such a document. The resulting scenarios resemble a more detailed form of the descriptions that are presented in Table 13.1. This identified the reporting requirements that must be satisfied by the various accident and incident reports that are published by the Hong Kong Marine Department. As the drafting of a report progresses, investigators can ask themselves whether the document that they have prepared might be used to complete the tasks that are identified in each of the scenarios. The 'evaluation' continues by showing a colleague what it would be like to use the document to achieve these identified tasks. This technique an be used at the very earliest stages of drafting a report and hence is a useful approach to formative evaluation. The problems with the use of scenarios are that it can focus designers' attention upon a small selection of tasks and users. For instance, a scenario might require that the Director of Marine Operations should be able to rapidly identify any regulations that were violated in the course of an incident. Such a scenario would not provide confidence that an engineer would be able to use the same document to rapidly identify detailed design recommendations. A further limitation is that it is difficult to derive empirical data from the use of scenario based techniques. Investigators may be able to convince their colleagues that a draft report satisfies the proposed requirements. This need not increase confidence that others will reach the same conclusions.

- *Experimental Techniques.* The main difference between the various approaches to evaluation is the degree to which investigators must constrain the reader's working environment. In experimental approaches, there is an attempt to introduce the empirical techniques of scientific disciplines. It is, therefore, important to identify the hypothesis that is to be tested. The next step is to devise an appropriate experimental method. Typically, this will involve focusing in upon a particular aspect of the many tasks that might eventually be supported by an incident report. For example, a safety-manager might be asked to reconstruct the events leading to an incident after having read a report for some specified amount of time. The reconstructions might then be examined to demonstrate that potential readers can more accurately recall these events in one version of the report than in another.

  In order to avoid any outside influences, tests will typically be conducted under laboratory conditions. Individuals are expected to read the draft report without the usual distractions of telephones, faxes, other readers etc. The experimenter must not directly interact with the reader in case they bias the results. The intention is to derive some measurable observations that can be analysed using statistical techniques. There are a number of limitations with the experimental approach to evaluation. For instance, by excluding distractions it is extremely likely that investigators will create a false environment. This means that readers may be able to draw inferences from a report more quickly and with greater accuracy than might otherwise be obtained within a noisy, complex working environment. These techniques are not useful if investigators only require formative evaluation for half-formed hypotheses. It is little use attempting to gain measurable results if you are uncertain what it is that you are looking for.

- *Cooperative evaluation techniques.* Laboratory based evaluation techniques are useful in the final stages of summative evaluation. In contrast, cooperative evaluation techniques (sometimes referred to as 'think-aloud' evaluation) are particularly useful during the formative stages of drafting an incident report. They are less clearly hypothesis driven and are an extremely good means of eliciting feedback on the initial drafts of a document. The approach is extremely simple. The experimenter sits with the reader while they work their way through a series of tasks with a potential report. This can occur in the reader's working context or in a quiet room away from the 'shop-floor'. The experimenter is free to talk to the reader but it is obviously important that they should not be too much of a distraction. If the reader requires help then the investigator should offer it and note down the context in which the problem arose. These requests for help represent the 'gaps' identified by Freeman [280]. Additional evidence or warrants may be necessary to support the claims that are made in an incident report. The main point about this exercise is that the reader should vocalise their thoughts as they work with the draft. This low cost technique is exceptionally good for providing rough and ready feedback. Readers can feel directly involved in the drafting of a final document. The limitations of cooperative evaluation are that it provides qualitative feedback and not the measurable results of empirical science. In other words, the process produces opinions and not numbers. Cooperative evaluation is extremely bad if investigators are unaware of the political and other presures that might bias a reader's responses.

- *Observational techniques.* There has been a sudden increase in interest in the use of observation techniques to help 'evaluate' a range of computer-based systems [751], of management structures [396] and of safety-critical working practices [120]. This has largely been in response to the growing realisation that the laboratory techniques of experimental psychology cannot easily be used to investigate the problems that individuals can experience in real-world settings. Ethnomethodology requires that a neutral observer should enter the users' working lives in an unobtrusive manner. They should not have any prior hypotheses and simply record what they see, although the recording process may itself bias results. This approach provides lots of useful feedback during an initial requirements analysis. In complex situations, it may be difficult to form hypotheses about readers' tasks until investigators have a clear understanding of the working problems that face their users. This is precisely the situation that affects many incident reporting systems; the individuals who run these applications often have very limited information about the ways in which others in their organisation, or in other organisations, are using the information that they gather [748]. There is a natural concern to be seen to meet existing regulatory requirements. Consequently it can be difficult to interpret an organisations' response when asked whether or not a particular report has guided their operating practices. Ethnography focuses not on what an organisation says that it does but on what key individuals and groups actually do in their everyday working lives. Unfortunately, this approach requires considerable skill and time. It is extremely difficult to enter a working context, observe working practices and yet not affect the users behaviour in any way. There have been some notable examples of this work. For example, Harris has shown how these techniques can be used to improve our understanding of midwife's behaviour in a range of safety-critical applications [306]. Her work also illustrates the potential drawbacks of the approach; her observations are grounded in several years of experience observing their working environment.

As mentioned, these techniques have been widely applied to evaluate the usability and utility of computer-based systems and of more general forms of documentation [686]. They have not been widely used to have validate the presentation of incident reports. It is for this reason that McGill conducted a series of initial investigations to determine whether some of these approaches might support such evaluations [530]. This work formed part of a wider study that was intended to determine whether the use of electronic presentation techniques support or hinder the presentation of incident reports. These wider findings are discussed in the closing sections of this chapter. For now it is sufficient to observe that the tasks were designed to determine how well readers could use the published report to: identify key events that occurred at the same time during an incident but that occurred in different areas of the system; discover the timing and location of key events; identify whether or not an individual was involved in particular events. His study used laboratory-based, cooperative evaluation techniques but he also derived a number of high-level performance measures during these tasks. Brevity prevents a complete analysis of McGill's results. It is, however, possible to illustrate some of the findings. For instance, readers were relatively proficient at using paper-based reports to find answers to factual information. Questions of the form 'write down the time that the Chief Officer left the mess room to return to the bridge' were answered in an average of 4 minutes and 28 seconds using a relatively small sample of only five readers. In contrast, tasks that involved the resolution of conflict or that related to causal hypotheses took far longer to resolve. For instance, none of the users were able to answer 'Officer A gave conflicting evidence as to the time at which he left G deck to to the mess room - write down the page reference where these conflicting statements are highlighted'. Mc Gill's work, like that of Snowdon [748], is suggestive rather than conclusive. It provides an indication of some of the problems that readers face when attempting to use existing incident reports to perform a particular set of tasks. The study makes a number of recommendations about how those problems might be addressed for an individual incident report and hence indicates how lab-based studies might be used to support both formative and summative evaluation.

These preliminary studies raise more questions than they answer. The subject groups are too small to provide results that can easily be generalised beyond the specific context of the particular evaluations. Previous studies have also been conducted away from the reader's working environment. Snowdon's study looked at safety managers' attitudes during the 'atypical events' of a trade conference. McGill's study took place within a University laboratory. Neither of these environments approaches the ecological validity that is a focus for the ethnographic techniques that are summarised in previous paragraphs.

McGill's work also identified significant problems in accounting for learning effects. In order to demonstrate that any changes to an incident report had addressed the problems identified in previous studies, he was forced to retest subjects. Unfortunately, the readers of the reports were able to use the knowledge gained with earlier versions of the report to perform some of subsequent tasks. He, therefore, used elaborate counter-balancing to ensure that some of the readers were presented first with the revised version of the report and others with the initial version. Unfortunately this leads to further problems where preliminary drafts provide greater support for some tasks that the subsequent versions. The complex nature of prose incident reports make it very difficult to be certain about which particular aspects of a document actually support particular user tasks.

The previous observation leads to further reservations. As mentioned, we know remarkably little about the diverse ways in which readers exploit incident reports. The same documents have been used to inform the drafting of recommendations that are intended to avoid future incidents, they have also been analysed as part of wider statistical summaries, they have been added to international databases that can be searched by interactive queries etc [413]. It is, therefore, difficult to identify appropriate tasks that might help to drive any experimental evaluation. Those studies that have been conducted have all commented on the difficult of 'pinning down' reporting organisations to a precise list of the tasks that these documents are intended to support.

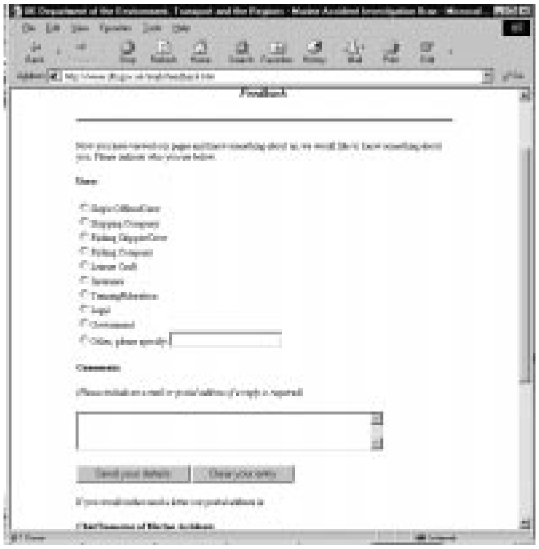

Feedback

[illegible]

○ Ship's Officer/Crew
○ Shipping Company
○ Fishing Skipper/Crew
○ Fishing Company
○ Leisure Craft
○ Insurance
○ Training/Education
○ Legal
○ Government
○ Other, please specify

Comments

[illegible]

Send your details | Clear your entry

[illegible]

Figure 13.7: MAIB On-Line Feedback Page

Even if we could agree upon an appropriate selection of tasks, it is difficult to envisage ways in which studies might move beyond the test-retest approach taken by McGill. This approach attempts to identify small differences between different formats by testing and retesting users with various versions of a report over a short period of time. It is likely, however, that any changes to the format of incident reports may be introduced as part of more systematic changes and that these changes could have consequences for the long-term application of the information that they contain. I am unaware of any longitudinal studies into the strengths and weaknessed of particular presentation styles for incident reports. This forms a strong contrast with the many studies that have been conducted into the usability and utility of design documents [94, 512]. None of the investigation agencies that I have contacted throughout the preparation of this book have reported the use of direct user-testing to improve the quality of their incident and accident reports. This does not imply that such organisations are not concerned to elicit feedback about the information that they provide. For instance, Figure 13.7 illustrates how the UK MAIB provides a web-based form to elicit feedback about its work. Unfortunately, there is no publically available information about the insights that such systems provide about the presentation of incident reports. It is instructive to observe, however, that this feedback page asks respondents to indicate whether their information related to "ship's officer/crew, shipping company, fishing skipper/crew, fishing company, leisure craft, insurance, training/education, legal, government or other". This diverse list again illustrates the problems of devising appropriate reports that might satisfy the different information requirements for all of these groups. The use of electronic feedback froms to elicit information from the intended readers of an incident report is a relatively new innovation. It reflects a far wider move towards the use of electronic systems in the presentation and dissemination of information about near-miss incidents and adverse occurrences.

13.4 Electronic Presentation Techniques

Many of today's accident reports exploit the same presentation format and structure that was first used at the beginning of the twentieth century. This would not be a concern if the nature of accident investigation had not changed radically over the same time period [248]. Chapter 3 has described how the introduction of computer technology has rapidly increased the integration of heterogeneous production processes. This has now reached the point where isolated operator errors or single equipment failures have the potential to create complex and long lasting knock-on effects within many different applications. Accident reports must now not only consider the immediate impact of an accident but also the way in which emergency and other support services responded in the aftermath of an incident. This was illustrated by the Allentown incident, described in Chapter 9. Technological change is only one of the factors that have changed the nature of accident reporting. There has, for example, been a move away from blaming the operator. This wider perspective has resulted in many regulatory organisations looking beyond catalytic events to look at the organisational issues that form the latent causes of many failures. These two factors, technological change and new approaches to primary and secondary investigation, have increased the length of many incident reports. The scale and complexity of such documents imposes clear burdens upon the engineers, managers and operators who must understand their implications for the maintenance, development and running of many safety-critical systems. Recently, however, a number of regulatory authorities have begun to use the web as a primary means of disseminating accident reports. The MAIB "are in the process of converting all reports into a format suitable for viewing and downloading from this web site" [518]. Electronic media are perceived to offer considerable support for the communication of these documents.

It is difficult to obtain direct evidence that the changing nature of incident investigation and the complexity of technological failures has increased the length of incident reports. Tables 13.8 and 13.9 show how the number of pages required by marine occurrence reports has risen between 1991 and 1999 for the ATSB. These two years were chosen because 1991 is the earliest year for which it is possible to obtain a relatively complete collection of reports. The fact that it can be difficult to obtain particular documents is itself worth noting. 1999 is the most recent year for which a complete

1991	
Report Number	Number of Pages
27	27
28	28
29	Unavailable
30	44
31	23
33	Unavailable
42	21
32	29
34	18
35	34
36	23
37	38
38	29

Table 13.8: ATSB's 1991 Marine Report Page Counts

1999	
Report Number	Number of Pages
143	40
144	27
145	23
146	35
147	29
148	37
149	40
150	42
151	35
152	42

Table 13.9: ATSB's 1999 Marine Report Page Counts

collection can be obtained, given that some investigations that were started towards the end of 2000 have not published their findings. It was possible to obtain copies of 11 of the 13 reports that were published in 1991. These extended to an average of 28.5 pages, the standard deviation was 7.6 and the total number of pages was 314. In 1999, the average had risen to 35 pages per report over 10 incidents with a standard deviation of 6.63. A certain degree of caution should be exercised over the interpretation of these figures. For instance, the style of presentation has changed radically over this period. In particular, the introduction of photographic evidence has considerably lengthened later reports. It can also be argued that this trend is atypical. The ATSB reports focus on high-risk incidents, they also represent the publications of a single national agency.

Further evidence can be found to support the hypothesis that changes in the scope and complexity of incident investigations has had the knock-on effect of increasing the length of many incident reports. Table 13.10 illustrates how word counts can be used to strip out the formatting differences that distort the page counts shown in Tables 13.8 and 13.9. These word counts were derived from the less critical incidents that are reported by the UK MAIB's Safety Digest publication. Unfortunately, a number of further problems affect such an analysis. The 1996 volume examined approximately 54 incidents while the 2000 edition considered more than 110. It is for this reason that Table 13.10 considers the word counts for the first five incidents in the first number of each volume for each of the years. As can be seen, however, the results of this analysis provide only partial support for

Volume and number	Year	No. of Words First 5 Incidents
Safety Digest 1/00	2000	2845
Safety Digest 2/99 1/99 unavailable	1999	2658
Safety Digest 1/98	1998	4577
Safety Digest 1/97	1997	2720
Safety Digest 1/96	1996	2178

Table 13.10: UK MAIB Safety Digest Incident Word Counts

the general hypothesis of increased page counts. The sudden rise in 1998 cannot be explained by a change in the investigatory process nor any sudden increase in the complexity of maritime incidents. In contrast, this sudden rise can be explained by the introduction of more detailed 'lessons learned' summaries into individual incident reports.

Recent initiatives to exploit electronic presentation and dissemination techniques cannot be explained simply in terms of the changing nature of incidents or new investigatory techniques. There are strong financial incentives that motivate the use of web-based systems to support incident reporting. There are considerable overheads involved in ensuring that investigatory and regulatory organisations have an 'up to date' inventory of previous reports. The growth of the Internet also offers the possibility of dissemination reports to organisations and individuals who might not have taken the trouble to order and pay for paper-based versions of an incident report. Of course, there are obvious security concerns associated with such electronic techniques. There are also a host of additional benefits for the automated indexing and retrieval of large scale incident collections. These retrieval issues will be discussed in the next chapter. In contrast, the following pages identify a range of techniques that are intended to support the electronic presentation of incident reports.

13.4.1 Limitations of Existing Approaches to Web-Based Reports

Two primary techniques have been used to support the electronic dissemination of incident reports: the Hypertext Mark-up Language (HTML) and Adobe's proprietary Portable Display Format (PDF). Figure 13.8 illustrates the use of HTML by the UK MAIB in their Safety Digest, mentioned above. Users can simply select hyperlinks to move between the pages that describe similar incidents.

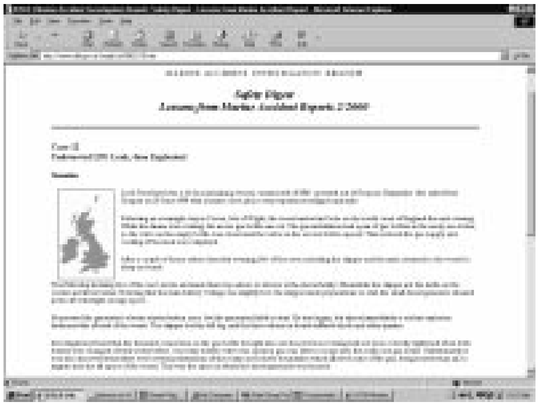

MARINE ACCIDENT INVESTIGATION BRANCH

Safety Digest

Lessons from Marine Accident Reports 2/2000

Figure 13.8: MAIB Safety Digest (HTML Version)

This use of the web offers a number of benefits for the presentation of incident reports. It avoids the overheads of maintaining a catalogue of paper-based documents. It is important to note, however, that very few agencies intend to entirely replace paper-based incident reports. For instance, the Transportation Safety Board of Canada include an explicit disclaimer on their web-site that:

> "These documents are the final versions of occurrence investigation reports as approved by the Transportation Safety Board. The TSB assumes no responsibility for any discrepancies that may have been transmitted with the electronic versions. The printed versions of the documents stand as the official record." [788].

HTML offers a number of further benefits. No special software is needed beyond a browser, such as Netscape Navigator or Internet Explorer, that are now installed as a standard feature of most personal computers. The introduction of hyperlinks and on-line keyword search facilities also helps to reduce the navigation problems that frustrate the readers of paper-based documents. There are, however, a number of limitations. Previous research has identified both perceptual and cognitive problems associated with the on-screen reading of technical documents [657]. This explains the subjective difficulties that reader's report when using large HTML documents. The improved comprehension that can be obtained through the appropriate use of hyperlinks as structuring tools can be jeopardised if on-line documents simply replicate the linear structure of paper based reports [670]. It can take almost twice as long to read electronic copy [555]. Readers are more prone to error when reading on-line documents [875]. HTML tags work well when they are interpreted and displayed by current web browsers. They do not work well when the same browsers are used to obtain printed output from HTML documents. Figures and photographs are often embedded as hypertext links in existing HTML reports. These will be missing in the printed version. Readers must manually piece together their hardcopy.

Many of these limitations can be avoided through the use of Adobe's proprietary PDF. It is for this reason that the MAIB publish their Safety Digest in both HTML and PDF formats. This exhaustive approach is, however, rare. Most agencies use either PDF or HTML. Figure 13.9 presents an excerpt from a PDF report into a marine incident published by the ATSB . The freely available PDF viewer integrates images and text to emulate the printed document on the screen. Readers can also obtain well-formatted, printed copies. This reduces the psychological and physiological problems of on-screen reading. However, these important benefits must be balanced against a number of problems. Firstly, it can be difficult for people to obtain and correctly install the most recent version of the PDF reader. This is important because these programmes are, typically, not a standard part of most browsers. Although PDF files are compressed, users can also experience significant delays in accessing these documents compared to HTML reports. Finally, it can be difficult to extract information once it has been encoded within a PDF document. This is a significant barrier if readers want to compile their own index of related incidents within an industry. The ATSB avoid this problem by providing extensive summaries of each incident in HTML format. Readers access the full PDF version of the report by clicking on a hyperlink within the HTML summary. This enables search and retrieval tools to index the HTML summary so that readers can easily find it using many of the existing search engines. They can then read and print the report in the PDF format that is less easy to index because of Adobe's proprietary encoding.

The hybrid use of both HTML and PDF by the ATSB and the MAIB does not address all of the problems that affect the electronic presentation of incident reports. These formats are, typically, used to reproduce the linear document structure that has been exploited since the beginning of the twentieth century. This imposes significant burdens upon the reader and may fail to exploit the full potential of web based technology [158]. It can be argued that investigatory authorities have focussed upon the electronic *dissemination* of accident reports over the web. Few, if any, have considered the opportunities that this medium provides for the effective *presentation* of these documents. The following paragraphs, therefore, describe how visualisation techniques from other areas of human-computer interaction can be used to support the electronic presentation of incident reports.

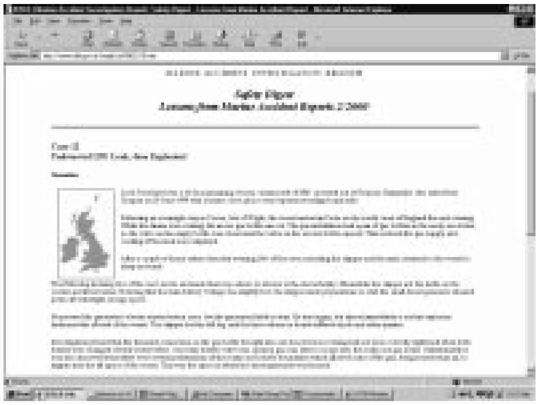

Figure 13.9: ATSB Incident Report (PDF Version)

13.4.2 Using Computer Simulations as an Interface to On-Line Accident Reports

Chapter 8 describes a range of simulation techniques that are intended to help investigators gain a better overview of the events leading to an incident. It is surprising that this approach is not more widely integrated into the on-line presentation of accident reports. For instance, Figure 13.10 shows how the NTSB's PDF report into a rail collision uses a still image from a 3-D simulation to provide readers with an overview of the train collision. This simulation was undoubtedly available to accident investigators. It would have been relatively simple to provide access to other readers alongside the PDF report. Instead the reader has to piece together events from more than 40 pages of prose.

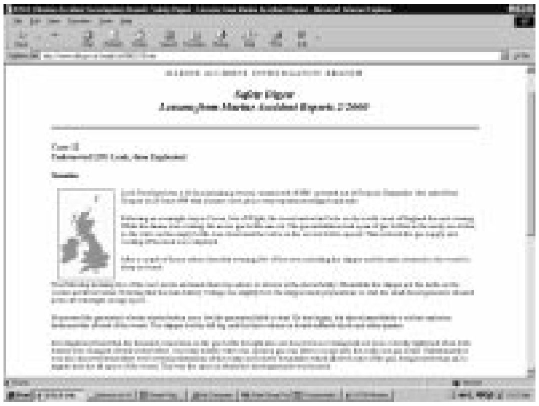

Figure 13.10: NTSB Incident Report (PDF Version)

Such examples, arguably, illustrate a missed opportunity to exploit novel means of presenting information about near-miss incidents and adverse occurrences. Several of my students have, therefore, begun to use simulations as an interface to on-line reports. Users are presented with animations of the events leading to an incident. Their browser also simultaneously present a set of links to those sections of the existing report that deal with the stage of the accident that is currently being simulated. The links are automatically updated as the simulation progresses. Users can stop a simulation at any point during its execution. They can then use their browser to retrieve the relevant sections of the text-based report.

Figures 13.11 and 13.12 illustrate two different applications of this approach. The system in

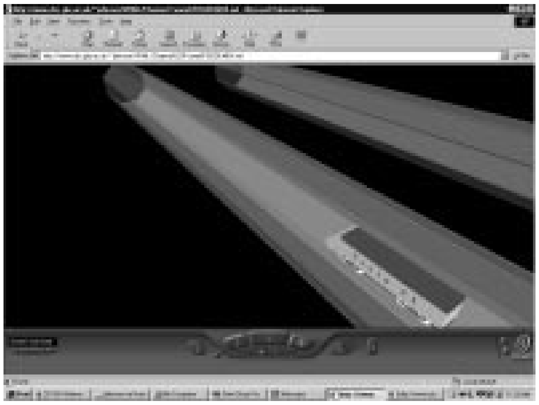

Figure 13.11: Douglas Melvin's Simulation Interface to Rail Incident Report (VRML Version)

Figure 13.11 uses the scripting facilities in the Virtual Reality Mark-up Language (VRML) to generate a scale model for part of the Channel Tunnel. The technical details of the VRML approach are briefly introduced in Chapter 8. A free helper application is integrated into the browser so that users can view the pseudo-3D images. The user can select areas of the image to replace the simulation with an HTML from the report into the incident. The intention is that the simulation will provide a hgh-level overview of the events leading to the failure. Readers can then use the animation as a means of indexing into the analysis and more detailed reconstructions that are presented in a more conventional format.

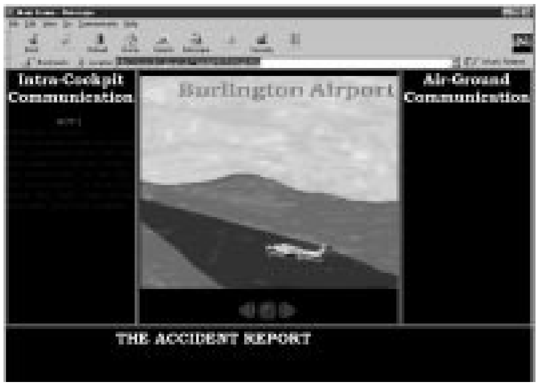

Figure 13.12: James Farrel's Simulation Interface to Aviation Incident Report (VRML Version)

The approach in Figure 13.12 exploits similar techniques but focuses on cockpit interaction during an air accident. A series of still images can be updated using the controls in the centre of the screen. On either side of the simulation are sections that present the transcripts both from the cockpit voice recorders and from the Air Traffic Controllers. The prose from the accident report is presented at the bottom of the screen. All sections of the interface are updated as the user moves through the simulation. Unfortunately, simulations cannot be used to support the presentation of all aspects of an incident report. Chapter 8 reviews the problems that arise when these techniques are used to model the distal causes of an incident. It is relatively easy to simulate the immediate events surrounding a particular occurrence, it is less easy to recreate the management processes and regulatory actions that create the context for an incident. Similarly, it can be difficult to represent near-misses or errors of omission. readers and 'viewers' often fail to detect that something which ought to have happened has not, in fact, been shown in the simulation.

13.4.3 Using Time-lines as an Interface to Accident Reports

The previous simulations attempted to recreate the events leading to failure. It is also possible to exploit more abstract visualisations to provide readers with a better overview of an incident [502]. For example, we have used the Fault tree syntax that was introduced in Chapter 9 as a gateway into an incident report. As shown in Figure 9.9, 9.10 and 9.12 these diagrams can be used to map out both the proximal and distal causes of an incident. They can also be annotated in various ways, for instance events can show the time or interval during which they are assumed to have occurred. Imagemap techniques provide a means of using these diagrams to directly index into an incident report. Imagemaps work by associating the coordinates of particular areas on an image with web-based resources. The net effect is that if the user clicks on a node in a Fault Tree the browser can be automatically updated to show those sections of prose in an incident report that are represented by the graphical component. However, such representations quickly suffer from problems of scale. This is a significant limitation if the scope and complexity of incident reports is increasing in the manner suggested by the opening paragraphs of this section.

Again the desktop virtual reality provided by VRML and similar languages can be exploited to ease some of these problems. We have developed the pseudo-3D time-line shown in Figure 8.7 and the perspective wall, shown in Figures 8.8 and 8.9, to provide novel means of interacting with incident reports. These visualisations enable readers to access reports over the web. Rather than having to scroll over large, two-dimensional imagemaps of graphical structures such as Fault Trees, users can 'walk' into these structures along the Z-plane. The claimed benefits of this approach include a heightened sense of perspective and the ability to focus on particular aspects of the structure by choosing an appropriate viewpoint. As before, readers can use their mouse to select areas of the 3-dimensional models. The system then automatically updates an area of the browser to present the relevant areas of a textual report.

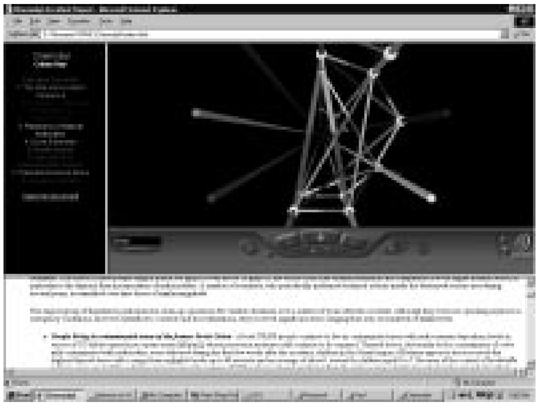

Figure 13.13: Peter Hamilton's Cross-Reference Visualisation (VRML Version)

Unfortunately, the models shown in Figure 8.7, 8.8 and 8.9 were all built manually . This requires considerable skill, expertise and patience. Together with Peter Hamilton, I have developed a similar range of three-dimensional interfaces that can be directly generated from textual incident report. This approach is illustrated in Figure 13.13. A colour-coded index is presented on the left of the image. This refers to each of the bars on the three dimensional representation shown on the right hand side. The image on the right represents time advancing into the Z plane. Each bar, and therefore each colour in the index, relates to a chapter in the incident report. Links are drawn between any two bars that refer to the same instant in time during the accident. It is, therefore, possible to 'walk' into the structure on the right to see whether different chapters treat the same events from slightly different perspectives. The lower area of the browser is used to present the HTML version of the report. Selecting a bar at any point in the Z plane will result in the image being updated to show any text relating to that moment in time in the selected chapter .

Many of these tools satisfy two different sets of requirements. For instance, the perspective wall and the three-dimensional time-line were introduced in Chapter 8 as tools to help investigators reconstruct the events leading to an incident. In this chapter, we have also argued that these techniques might also support the presentation of accident and incident reports. Similar comments can be made about the use of location-based simulations. For example, the QuicktimeVR images shown in Figure 8.4 have been integrated into several incident reports. Similarly, the imagemap shown in Figures 8.1 and 8.2 was initially intended as an alternative interface to an incident report rather than as an aid to incident investigators .

These interfaces also support a number of additional tasks. For instance, Schofield has pioneered the use of similar simulations within a range of legal proceedings [730]. This application of forensic animation is still very much in its infancy. The Civil Procedure Rules (CPR) substantially changed the legislative framework that governed the use of these systems in English courts. These rules came into force in April 1999 and gave judges considerable powers to control both the conduct and preparation of trials [292]. However, there is insufficient precedent for the use of these techniques and so much of the focus of Schofield's work has been to build up collections of 'case law' from other legal jurisdictions.

It is clear that the cost associated with the more innovative visualisations proposed in this section can only be justified for high-risk incidents. On the other hand, all of the interfaces shown here were produced using mass-market, software that is available without charge or under shareware license agreements. It is also important to stress that the development costs were minimal. They were developed over a period of approximately 1-2 weeks by final-year undergraduate students participating in my course on safety-critical systems. The intention behind this exercise was to see whether people who knew little about incident reporting but who knew more about the computer-based presentation of information might come up with some innovative presentation ideas.

There are, however, a number of caveats that must be raised about the potential benefits of these innovative presentation techniques. These can be summarised as follows:

1. *The problems of navigating in current desktopVR environments.* Many users experience considerable difficulty in moving along the three dimensional time-lines shown in the previous paragraphs. This problem has been widely reported in previous work on the validation of desktopVR [202]. Several solutions have been proposed. For instance, recent versions of our software have made extensive use both of virtual way-points and of sign posts [637]. This approach takes users on a 'conducted tour' of an incident . Users do not have to concentrate on achieving a particular orientation at a particular point in three-dimensional space in order to view particular events. They can simply click on a menu item that will take them to those events. Their viewpoint and orientation in the virtual world is automatically updated to provide them with the best view of the surrounding events. These navigation problems do not affect all accident simulations. For instance, the interface in Figure 13.12 uses animation techniques that enable the developer to treat the simulation as a 'movie'. They can specify the viewpoint and the sequence of events so that the reader need not navigate within a three dimensional environment.

2. *Unknown rhetorical effects.* Many of the techniques that are presented in this paper introduce new rhetorical techniques, or tropes. The introduction of a simulation can have a profound impact upon many readers. Initial field trials have indicated that people are more willing to believe the version of event shown in a simulation or model than they are in a paper based r eport. On the one hand this illustrates the importance of these new techniques. Investigation agencies can use them as powerful tools to convince readers about a particular view of an incident. On the other hand, these new techniques may persuade people to accept a simulated version of events that cannot be grounded in the complex evidence that characterises modern accidents.

3. *The problems of validating novel interfaces.* DesktopVR interfaces and simulations have a subjective appeal [406]. This should not be under-estimated. Investigatory and regulatory agencies have organisation reasons for being seen to embrace new technology. The Rand study

> [482] and the Institute of Medicine report [453] cited in previous chapter both criticised a range of perceived problems that have limited the effectiveness of more 'traditional' approaches to incident reporting. This pressure to innovate can be reinforced by the strong subjective appeal, mentioned above. However, the attraction can quickly wane. There is a danger that the users of these novel interfaces will reject the additional facilities that they offer if those techniques are not perceived to support their diverse working tasks. Previous sections have described the very limited nature of those studies that have been conducted into the longitudinal use of conventional, paper-based incident reports. We have even less information about the potential long-term use of these more innovative approaches.

There have been some limited evaluations of these novel presentation techniques. For instance, the McGill study mentioned in previous sections was extended to compare the use of a VRML time-line with the use of a location-based imagemap. The imagemap presentation was introduced in Chapter 8 and is illustrated in Figures 8.1 and 8.2. As mentioned, the readers of the report use a web browser to view a cross-sectional diagram of the ship that was involved in the incident. By selecting areas of that image, they can access a time-line of the events that happened in particular locations. These events are accompanied by a brief summary of their importance within the wider causes of an incident.

McGill's laboratory evaluation also considered the use of the VRML time-line illustrated in Figure 8.7 as a presentation format for incident reports. Readers used a web-based interface to 'walk' or 'fly' through an abstract map of the events leading to a particular adverse occurrence or near-miss incident. By clicking their mouse on certain areas of the structure, they accessed web pages that presented more detailed textual or graphical information. The participants in the study were asked to perform a number of tasks within a fixed period of time .

In order to minimise the effects of fatigue, each participant was stopped when they had spent a maximum of five minutes on any individual task. The study was counter-balanced so that each task was performed using a different presentation technique and the same number of participants performed each task using a particular approach. McGill also altered the order in which particular tasks were performed to help minimise any effects of fatigue or of learning from answering previous questions using a different presentation technique. The five tasks were as follows:

1. Write down the time that the Chief Officer left the mess room to return to the bridge.
2. Write down the key events which happened on the bridge at approximately 18:23.
3. Officer A gave conflicting evidence as to the time at which he left deck G to go to the mess. Write down the page/paragraph reference where these conflicting statements are highlighted.
4. What events happened at 18:28?
5. Write down the time that crewmember B returned to his cabin.

Table 13.11 summarises the times that were obtained for five users performing these different tasks using cooperative evaluation under laboratory conditions. The small number of participants in the study meant that a very limited number of readers used each interface to perform each task. It can also be argued that his choice of questions may have produced results that are favourable towards particular presentation formats. McGill concludes that the electronic versions offered significant benefits over the paper-based presentation of the incident report. This is revealed principally in the time take to perform the specified tasks but also in a range of attitudinal questions that assessed the readers' subjective response to these systems. These subjective questions indicated a strong preference for the VRML time-line over the image map for most tasks. McGill did not compare subjective responses for the paper-based report because he assumed that the electronic formats would supplement rather than replace more traditional presentation techniques.

The results summarised in Table 13.11 hide a number of factors that influenced the course of the study. For instance, subject 3 in Task 4 missed one of the events. McGill awarded a time penalty for each incorrect answer when presenting his results [530]. His decision to add thirty seconds to the time reported for task completion seems arbitrary. These penalties have not, therefore, been

	Paper report	Image Map	VRML Time-line
User 1	Task 1: Did not complete Task 4: 133 sec.	Task 2: 52 sec. Task 5: 144 sec.	Task 3: 35 sec.
User 2	Task 2: 227 sec. Task 5: 173 sec.	Task 3: 62 sec.	Task 1: 19 sec. Task 4: 45 sec.
User 3	Task 3: Did not complete	Task 1: 189 sec. Task 4: 179 sec.	Task 2: 38 sec. Task 5: 27 sec.
User 4	Task 2: Did not complete	Task 1: 177 sec. Task 5: Did not complete	Task 3: 61 sec. Task 4: 69 sec.
User 5	Task 1: 236 sec. Task 5: 254 sec.	Task 2: 40 sec. Task 4: 273 sec.	Task 3: 42 sec.

Table 13.11: McGill's Timing Results for Tasks with Electronic Incident Reports

included in Table 13.11. The preliminary results provided by this study must be supported by more sustained investigations into the potential errors that might arise from the use of such innovative presentation techniques. A number of further comments can be made about these results. For example, user 4 in task 5 insisted on abandoning the image map and resorted to the paper based report after 3 minutes. Such situations typify experimental studies with users performing complex tasks over even relatively short periods of time. There are inevitable frustrations from being forced to use presentation formats that the reader might not accept if they were given a free choice. This emphasises the need for more sustained longitudinal studies into the 'real world' use of electronic incident reports.

13.5 Summary

Incident reports provide a primary means of ensuring that the weaknesses of previous applications are not propagated into future systems. This chapter has described a number of the problems that complicate the task of drafting these documents. Incident reports may have to be tailored to meet the particular requirements of several diverse audiences. For instance, some reports are produced to be read by the regulators who must decide whether or not to act on their recommendations. Other forms of incident report are intended as case studies or scenarios that are to be read by operators and practitioners. Further factors complicate even these relatively simple distinctions. For instance, the format and presentation of reports that are to be read by operators within the same organisation as the investigator can be quite different from those that may be disseminated to a wider audience. The problems of addressing the needs of intended readers of an incident report are also complicated by the way in which different types od report may be drafted for different types of occurrences. For example, incidents that have a high-risk associated with any future recurrence may warrant greater detail than those that are assumed not to pose any future threat to a system. Different presentation techniques may also be necessary in order to draft reports at different stages of an investigation. The format of an interim report in the immediate aftermath of a near-miss incident or adverse occurrence is unlikely to support the more polished, final report that is required by international organisations, such as the IMO. The task of meeting these various requirements is also exacerbated by the need to preserve confidentiality in the face of growing media and public interest in technological failures.

Subsequent sections of this chapter have presented a number of detailed recommendations that are intended to support the presentation of incident reports. These recommendations are ordered in terms of the different sections that are, typically, used to structure more formal reports into high-risk incidents: reconstruction then analysis then recommendations. However, many of the detailed guidelines can also be applied to documents that analyse less critical failures. For example, any reconstruction should consider both the distal and the proximal events that contributed to an incident. It is also important to describe the evidence that supports a particular version of events. Similarly, the analysis section should not only describe the causes of an incident but also the methods that were used to identify those causes. It is no longer sufficient in many industries to rely simply on the reputation of domain experts without the additional assurance of documented methods to support their conclusions [280]. The guidelines for drafting recommendations include a requirement that investigators should define conformance criteria so that regulators can determine whether or not particular changes have satisfied those recommendations. It is also important to explain the relationship between any proposed changes and the causes that were identified from any analysis of the incident.

It is important that investigators have some means of determining whether or not a particular report supports the tasks that must be performed by its intended audience. This can be done in one of two ways. Firstly, we have described a range of verification techniques that can be used to demonstrate that particular documents satisfy principles or properties that are based on the guidelines, summarised in the previous paragraph. Careful reading of an incident report can also be used to identify a range of rhetorical techniques. These tropes are often effective in more general forms of communication but can lead to ambiguity and inconsistency within engineering documents. Such problems can have profound consequences for safety when they weaken the presentation of incident reports. Mathematical proof techniques provide one means of identifying where enthymemes weaken the use of syllogism in an incident report. This is particularly important because the ommision of evidence not only removes particular information from the reader but it also prevents them from drawing necessary inferences that depend upon the missing information. We have also described how Toulin's [775] model was produced as a response to criticisms of logic as a means of analysing argument. This approach rejects any a priori notions of truth or falsehood. Instead it links evidence and claims through the warrants and backing that support them. It also emphasises the importance of qualifiers and rebuttals in explaining why someone might hold a contrary opinion.

We have identified limitations that affect both the use of logic and of models of argumentation to support the verification of incident reports. Neither approach provides any guarantee that the properties which are established of any particular document will actually be significant to their intended readership. Subsequent sections, therefore, briefly described a range of user-testing techniques that can support the validation of particular formats. Surprisingly few of these techniques have ever been applied to determine whether incident reports actually support the activities of their end-users. The results of some preliminary studies have been described, however, these raise a host of methodological problems that complicate such validation activities. For example, it can be difficult to identify a representative set of tasks that must be supported by a particular incident report.

This lack of evidence about the utility of existing formats is worrying given that a number of factors are creating potential problems for their continued application. Existing paper-based reporting techniques have, however, remained relatively unchanged over the last century. During this time, the introduction of microprocessor technology has significantly increased the integration of heterogeneous processes. There have also been changes to the techniques that drive incident analysis. There is an increasing awareness that contextual and organisational issues must be considered in addition to any narrow findings about operator 'error'. The integration of safety-critical interfaces and the development of organisational approaches to incident analysis have created new challenges for the presentation of these reports. Their scope and complexity is increasing. At the same time, many investigatory organisations have begun to disseminate their incident reports over the World Wide Web. It is, therefore, possible to extend the application of advance visualisation and navigation techniques from other areas of information technology to support the readers of these documents. Unfortunately, most agencies focus on the Web as a means of dissemination rather than communication. Paper based reports are directly translated into HTML or Adobe's PDF. This creates a number of problems. The HTML approach suffers from the well-known problems of reading large on-line documents. The PDF approach prevents many search engines from effectively retrieving information about similar accidents. Finally, the direct translation of paper-based reports into on-line documents ignores the potential flexibility of electronic media. The closing sections of this chapter have described a number of alternative approaches that support the presentation of incident reports. These formats exploit the opportunities created by the electronic dissemination and communication of incident reports. They must, however, be treated as prototypes. Further evidence is required to demonstrate that they support the range of tasks that are currently achieved using paper-based documents and their more orthodox counterparts on the web.

The increasing use of electronic media to document information about adverse occurrences and near-miss incidents not only creates opportunities for new presentation formats. It also, for the first time, creates opportunities to provide automated tools that will help investigators find records of similar incidents in other organisations, in different industries, in other countries. It is far from certain that we will be able to exploit this opportunity. For example, the vast number of records that have been compiled in some industries now makes it practically impossible to accurately search and retrieve information about similar mishaps. The following chapter, therefore, describes a range of computer-based tools and techniques that have been specifically developed to support large-scale collections of incident reports.

Chapter 14

Dissemination

The previous chapter looked at the problems associated with the presentation of incident reports. It was argued that the format and structure of these documents must be tailored so that they support their intended recipients. It was also argued that care must be taken to ensure that the rhetoric is not used to mask potential bias in an incident report. This chapter goes on to examine the problems that are associated with the dissemination of these documents. It is of little benefit ensuring that reports meet presentation guidelines if their intended recipients cannot access the information that they provide. There are significant problems associated with such dissemination activities. For example, the FDA's Medical Bulletin that presents information about their MedWatch program is currently distributed to 1.2 million health professionals. Later sections analyse the ways in which many organisations are using electronic media to support the dissemination of incident reports. This approach offers many advantages. In particular, the development of the Internet and Web-based tools ensures that information can be rapidly transmitted to potential readers across the globe. There are, however, numerous problems. It can be difficult to ensure the security and integrity of information that is distributed in this way. It can also be difficult to help investigators search through the many thousands of incidents that are currently being collected in many of these systems. The closing sections of this chapter present a range of techniques that are intended to address these potential problems.

14.1 Problems of Dissemination

Chapters 12 and 13 have already described some of the problems that complicate the dissemination of information about adverse occurrences and near miss incidents. For example, it can be difficult to ensure that information is made available in a prompt and timely fashion so that potential recurrences are avoided. It can also be difficult to ensure that safety recommendations reach all of the many different groups that might make use of this information. The following pages build on these previous chapters to analyse these barriers to dissemination in greater detail.

14.1.1 Number and Range of Reports Published

It is important not to underestimate the scale of the task that can be involved in the dissemination of incident reports. Even relatively small, local systems can generate significant amounts of information. For instance, one of the Intensive care Units that we have studied generated a total of 111 recommendations between August 1995 and November 1998. 82 of these were 'Remind Staff' statements. The 29 other recommendations concerned the creation of new procedures or changes to existing protocols (e.g. 'produce guidelines for care of arterial lines'), or were equipment related (e.g. 'Obtain spare helium cylinder for aortic pump to be kept in ICU').

As one might expect the task of keeping staff and management informed of recent incidents and recommendations is significantly more complex in national and international systems. This is illustrated by Table 14.1, which presents the total number of different reports that were published by

	1995	1996	1997	1998	1999	2000	2001 (-Aug)
Hazard Notices	6	12	16	2	6	13	2
Device Bulletins	5	7	4	6	3	5	4
Safety Notices	33	40	20	43	41	28	24
Device Alerts	-	-	-	-	-	8	5
Advice Notices	-	-	-	-	6	1	0
Pacemaker Notes	6	6	3	4	7	4	4
Total Reports	50	65	43	55	63	59	39
Total Incidents	4,298	4,330	5,852	6,125	6,860	7,352	-

Table 14.1: Annual Frequency of Publications by the UK MDA

the UK Medical Devices Agency (MDA) over the last five years. It should be noted that the figures for 2001 are currently only available until August. As we have seen from the Maritime examples in the previous chapter, incident reporting agencies produce a range of different publications to disseminate their recommendations. In Table 14.1, hazard notices are published following death or serious injury or where death or serious injury might have occurred [542]. A medical device must also be clearly implicated and immediate action must be necessary to prevent recurrence. Device bulletins address more general management interests. They are derived from adverse incident investigations and consultation with manufacturers or users. They are also informed by device evaluations. In contrast, safety notices are triggered by less 'serious' incidents. Their are published in circumstances where the recipients' actions can improve safety or where it is necessary to repeat warnings about previous problems. Device alerts are issued if there is the potential for death or serious injury particularly through the long term use of a medical device. Finally, Pacemaker Technical Notes publish information about implantable pacemakers, defibrillators and their associated accessories. For the purpose of comparison, Table 14.1 also contains the total number of adverse incidents that were reported in the MDA's annual reports [539]. As can be seen, there has been a gradual rise in the frequency of incident reports while the total number of publications has remained relatively stable. Such an analysis must, however, be qualified. The total incident frequencies are based on the MDA's reporting year. Hence the figure cited for 1996 is, in fact, that given for 1996-1997. However, the number of reports and associated publications provides some indication of the scale of the publication tasks that confront organisations such as the MDA.

	1997	1998	1999	2000	2001 (-Aug)
Safety Alerts	55	55	54	67	38
Drug Labeling	239	519	512	505	241
Biologics Safety	-	24	10	29	14
Food and Applied Nutrition	3	2	2	2	2
Devices and Radiology	9	12	9	4	3

Table 14.2: Annual Frequency of Publications by the US FDA's MedWatch Program

The level of activity indicated in Table 14.1 is mirrored by the figures in Table 14.2. This presents publication figures for the US Food and Drug Administration's MedWatch initiative. This Safety Information and Adverse Event Reporting Programme is intended to 'serves both healthcare professionals and the medical product-using public' [269]. It covers a braid range of medical products, 'including prescription and over-the-counter drugs, biologics, dietary supplements, and medical devices'. It, therefore, has a slightly wider remit than that of the UK MDA. The primary MedWatch publication provides Safety Alerts about drugs, biologics, devices and dietary supplements. As can be seen in Table 14.2, the MedWatch programme also publishes information from several different groups within the FDA. It publishes safety-related drug labeling change summaries that have been approved by FDA Center for Drug Evaluation and Research. It also incorporates recalls, withdrawals and safety issues identified by the Center for Biologics Evaluation and Research. The program also publishes selected warnings and other safety information identified by the Center for Food Safety and Applied Nutrition. Finally, the Medwatch initiative incorporates safety alerts, public health advisories and notices from the Center for Devices and Radiological Health. We have not calculated totals for Table 14.2 as we did for Table 14.1 because of the inherent difficulty of calculating the frequency of recommendations to change drug labelling in the FDA adverse event reporting programme. Some drugs form the focus of several reports in the same year. Recommendations can be applied to a particular generic named product or to the different brands of that product. We have chosen to calculate frequencies on the basis of named drugs identified in the Center for Drug Evaluation and Research warnings.

It is also important to stress that the reports identified in Tables 14.1 and 14.2 only represent a small subset of the publications that the FDA and the MDA publish in response to adverse incidents. For instance, the MedWatch programme also disseminates articles that are intended to support the continuing education of Healthcare professionals. These include information about the post-marketing Surveillance for Adverse Events After Vaccination and techniques for assuring drug quality. The FDA also provides more consumer oriented publications to encourage contributions from the general public. It uses incident information to address specific consumer concerns, for instance in special reports on drug development and the interaction between food and drug. The wide scope of these publication activities is also illustrated by the User Facility Reporting Bulletins. This quarterly publication is specifically targeted at hospitals, nursing homes and 'device user facilities'.

Dissemination activities are not only focussed on the generation of specific incident reports or articles on more general issues. They also include the organisation of workshops, such as the 1998 meeting on 'Minimising Medical Product Errors' [261]. The UK MDA host similar events, such as their annual conference which in 2001 will address the theme 'Protecting Patients - Maintaining Standards' [545]. The MDA also holds more focussed study days. These provide staff training to address common problems with particular devices. For example, the MDA set up a recent meeting for nurses on best practice and the potential pitfalls in operating infusion systems [543].

14.1.2 Tight Deadlines and Limited Resources

The previous paragraphs illustrate the high frequency and the diversity of dissemination activities that are conduced by many incident reporting organisations. It is difficult to under-estimate the logistical challenges that such activities can impose upon finite resources. There is also increasing pressure in many sectors to increase the efficiency of many reporting bodies. For instance, one government measure estimates that the MDA managed to increase its output by 9% with a stable workforce between 2000-2001. These pressures can also be illustrated by some of the objectives being promoted by the MDA. For 2001-2002, it is intended that all Hazard Notices will be issued within 20 working days; 90% of Safety Notices will be issued within 60 days and 75% within 50 days. It is also intended to increase the number of adverse incident reports that will be published by a further 9% while at the same time making 'efficiency' savings of 2% [539].

The results of tight financial constraints can also be seen in the manner in which the FDA has altered it's publication policy in recent years [867]. Previous sections have mentioned the User Facility Reporting Bulletin, this publication is intended for hospitals, nursing homes and other end-user facilities. The initial twenty, quarterly issues of the Bulletin were printed in the conventional manner and were posted to any organisation that requested a copy. At its peak, 77,000 subscribers received copies of these documents that presented summarised reports based on recent incident reports. Budgetary restrictions forced the FDA to review this policy. In Issue 17, readers were asked to respond to a retention notice. If they did not respond then they were removed from the distribution list. It was hoped that the high initial administrative overhead associated with this initiative would yield longer term savings in distribution costs. By 1999, however, Federal funding cuts prevented any distribution in paper form. The twenty-first issue of the Bulletin was, therefore, distributed through electronic means including an automated Fax system. The FDA summarised their feelings about this situation; 'we regret the need to move to this new technology if it means that many of our current readers will no longer have access to the Bulletin' [867].

The joint pressures imposed by the need to disseminate safety information in a timely fashion and the need to meet tight financial objectives has resulted in a number of innovations in incident reporting. Many of these systems start with the premise that it is impossible to elicit and analyse voluntary incident reports across an entire industry. Even with mandatory reporting systems there will be problems of contribution bias that result in a very partial view of safety-related incidents. These problems stem partly from the cost and complexity of large scale voluntary reporting systems. They also stem from the passive nature of most mandatory systems that simply expect contributors to meet their regulatory requirements when responding to an adverse occurrence. As we have seen, even if a potential contributor wants to meet a reporting obligation they may fail to recognise that a safety-related event has occurred. Most regulatory and investigatory organisations lack the resources necessary to train personnel across an industry to distinguish accurately between reportable and non-reportable events. Similarly, there are significant financial barriers that prevent routine inspections and audits to review compliance.

Sentinel reporting systems provide an alternative solution that is intended to reduce the costs associated with incident reporting and, thereby ensure that recommendations are disseminated in a timely fashion. This approach identifies a sample of all of the facilities to be monitored. This sites within the sample are then offered specialist training in both mandatory and voluntary incident reporting. The incidents that are reported by these sentinel sites can then form a focus for more general safety initiatives across an industry. These ideas are extremely suasive to many governmental organisations. For instance, the FDA Modernisation Act (1997) required that the FDA make a report to Congress in late 1999 about progress towards such a sentinel system [262]. In September 1996, CODA Inc. was awarded a contract to conduct a study to evaluate the feasibility and effectiveness of a sentinel reporting system for adverse event reporting of medical device use. The explicit intention was to determine not whether a sentinel system could supplement passive, voluntary systems, such as MedWatch, but to replace them entirely. The trial ran for twelve months and the final report emphasised the importance of feedback and dissemination in the success of any sentinel system. The CODA trial provide several different forms of feedback. These included a newsletter, faxes of safety notices, responses to questions presented by Study Coordinators. The individual reports that were received by the project were summarised, anonymised and then published in bimonthly newsletters for Study Coordinators. These coordinators acted as an efficient means of disseminating safety related information within the sample sites.

This project not has important implications for the efficient dissemination of safety-related information. It also provides important insights about the practical problems that can arise when attempting to ensure the timely dissemination of incident recommendations and reports. Many reporting systems endeavour to ensure that operators, safety managers and regulators are provided with information about incidents according to a sliding timescale that reflects the perceived seriousness of the incident. This is the case with the MDA targets, cited above. It can, however, be extremely difficult to estimate the seriousness of an incident. Previous chapters have referred to the 'worst plausible outcome' that is often invoked to support such assessments. The practical problems of applying such heuristics can be illustrated by the CODA pilot study. The project analysts determined that only 14% of the reports received would have been clearly covered by the existing mandatory systems. 56% of the reports described less serious incidents that fell within the voluntary reporting provisions. 30% of all submission, or 96 reports, fell between these two categories; 'the determination of serious patient injury according to FDA's definition was difficult to make'. Of these 96, 60% were submitted on voluntary forms. 25% of the reports clearly documenting serious patient injury also were submitted on voluntary forms. If these results provide an accurate impression of the true severity of the incidents then they indicate that analysts cannot accept the contributors' severity assessments at face value. Two senior nurse-analysts agreed to review all reports and classify them urgency using a scale of: very urgent, urgent, routine monitoring, well-known problem or not important. Approximately one-third (113) were classified as very urgent or urgent. Of these, only 19 were clearly mandatory reports. This is a significant concern given that distribution deadlines focus on a rapid response to mandatory reports.

The results of this analysis can be presented in another way. As mentioned, 14% of all reports clearly fell within the existing mandatory systems. About half of these, according to the nurses' analysis, needed only routine monitoring. The FDA cite the example of a 'problem with a catheter in which there was medical intervention, but for which FDA already had taken action, so that additional reports would not make a very valuable contribution to the agency' [262]. However, 50 of the 175 reports that fell under voluntary reporting rules were rated as very urgent or urgent. This creates considerable problems for the prompt dissemination of safety-related information. Delays in a regulatory response do not simply stem from the time required to analyse a report and make recommendations. They also stem from the amount of time that it takes a contributor to actually gfile a report in the first place. Some of the contributors complained about the time limits that were recommended for reporting particular classes of incidents. In some cases, what contributors classed as less severe occurrences went unreported for more than ten days. The previous paragraphs have questioned the reliability of such severity assessments and so it seems likely that such delays may be a significant factor in ensuring the prompt dissemination of alerts and warnings.

14.1.3 Reaching the Intended Readership

Chapter 13 argued that the task of drafting incident reports is complicated by the diverse readerships that these documents can attract. This section extends this argument. The diverse readership of these documents not only complicates the drafting of an incident report but also exacerbates their dissemination. There is an immediate need to ensure that individuals within a working unit are informed of any recommendations following an incident. Chapter 5 argued that such actions are essential to demonstrate that contributions are being acted on in a prompt manner. In particular, reports should be sent to the individuals who initially provided notification of an adverse occurrence. These requirement apply to the dissemination of incident reports in both large and small scale systems. National and international schemes face additional distribution problems. In particular, incident reports must forwarded to other 'at risk' centres. This is a non-trivial requirement because it can often be difficult to determine precisely which centres might be affected by any potential recurrence. Within these associated working groups, it is important to identify who will assume responsibility for ensuring the reports are read by individual members of staff. This can involve close liaison between the investigators who draft a report and safety managers or other senior staff distributed throughout their organisation.

It is possible to identify a number of different dimensions that characterise the distribution of incident reports. The following list summarises these different dimensions. Particular reporting system may tailor the approach that they adopt according to the nature of the incident. They may also use hybrid combinations of these techniques. For example, a closed distribution policy might be exploited within the organisation that generate the report to ensure that information was not prematurely leaked to the media. However, a horizontal approach might also be used to ensure that key individuals in other companies are also made aware of a potential problem:

- *Closed distribution.*
  This approach restricts the dissemination of incident reports to a few named individuals within an organisation. This creates considerable problems in ensuring that those individuals and only those individuals actually receive copies of a report. It is also important to note throughout this analysis that the receipt of a report does not imply that it will be either read or acted upon.

- *Horizontal distribution.*
  This approach allows the dissemination of incident reports to other companies in the same industry. The distribution may be further targeted to those organisations that operate similar application processes.

- *Vertical distribution.*
  This approach allows the dissemination of reports to companies that occur within the same supply chain as the organisation that was notified about an incident. Reports can be passed

down the supply chain to ensure that companies, which rely on the products and services of the contributor organisation, are altered to a potential problem. Supply companies may also be informed if an incident occurs as the result of problems at previous stages in the supply chain.

- *Parallel distribution.*
  This approach ensures that reports are distributed to companies in *other* industries that operate similar processes. For example, incidents involving the handling and preparation of nuclear materials can have implications in the defence, medical and power generation industries. It is for this reason that organisations such as the US Chemical Safety and Hazard Investigation Board were set up to span several related domains.

- *Open distribution.*
  This approach allows the free distribution of incident reports. Increasingly, this approach is being adopted by regulatory organisations, including the FDA and MDA, and by independent research organisations, such as the NHS Centre for Reviews and Dissemination [191] As we shall see, these open publication initiatives increasingly rely upon Internet-based distribution techniques.

The healthcare industry provides extreme examples of the problem associated with distributing incident reports to a diverse audience. In 2000, the MDA received 7,249 reports of adverse incidents involving medical devices. These resulted in 4,466 investigations after an initial risk assessment. 49 safety warnings were published; the MDA's annual report bases this figure on the sum of the numbers of Hazard Notices, Safety Notices and Device Alerts in Table 14.1 [539]. Safety Notices are primarily distributed through the Chief Executives of Health Authorities, NHS Trusts and Primary Care Trusts as well as the directors of Social Services in England. These individuals a responsible for ensuring that they are brought to "the attention of all who need to know or be aware of it" [535]. Each Trust appoints a liaison officer who ensures that notices are distributed to the 'relevant managers'. Similarly, each local Health Authority appoints a liaison officer to ensure that notices are distributed to Chairs of Primary Care Groups, Registration Inspection Units, Independent Healthcare Sector and representatives of the Armed Services. The MDA also requires that notices are sent to the Chief Executives of Primary Care Trusts who are then responsible for onward distribution to their staff. Social Services Liaison Officers play a similar role but are specifically requested to ensure distribution to Registration Inspection Units and Residential Care Homes.

The distribution responsibilities of the individuals in the MDA hierarchy are presented in Table 14.3. The detailed responsibilities of each individual and group are, however, less important than the logistic challenges that must be addressed by the MDA when they issue a Safety Notice. For instance, any individual warning will only be sent to some portion of the total potential audience. Many Safety Notices are not relevant to the work of Social Services. In consequence, each published warning comes with a list of intended recipients. These are identified by the first level in the distribution hierarchy: Health Authorities, NHS Trusts, Primary Care Trusts and Social Services. Liaison officers are then responsible for ensuring that information is directed to those at the next level in the hierarchy. This selective distribution mechanism creates potential problems if, for example, a Social Service department fails to identify that a particular Safety Notice is relevant to their operations. The MDA, therefore, issue a quarterly checklist that is intended to help liaison officers ensure that they have received and recognised all applicable warnings.

The MDA distribution hierarchy illustrates a number of important issues that affect all reporting systems. There is a tension between the need to ensure that anyone with a potential interest in a Safety Notices receives a copy of the warning. This implies that Liaison Officers should err on the side of caution and disseminate them as widely as possible. On the other hand, this may result in a large number of potentially irrelevant documents being passed to personnel. The salience of any subsequent report might then be reduced by the need to filter these less relevant warning. These arguments, together with the expense associated with many forms of paper-based distribution, implies that Liaison Officers should target any distribution as tightly as possible. Later sections of this chapter will return to this tension when examining the generic problems of precision and recall in information retrieval systems.

Organisation	Liaison Officer forwards to	For onward distribution to
NHS Trust	Appropriate Manager	Relevant staff to include Medical Directors, Nurse Executive Directors, Directors of Anaesthetics, Directors of Midwifery, Special Care Baby Units/Pediatric Intensive Care, Maternity Wards, Operating Theatres, Ambulance NHS Trusts and Accident and Emergency Units.
Health Authority	Primary Care	Directors of Primary Care Local Representatives Committees Chief Executives of Primary Care Groups Individual GP Practices Dentists Opticians Pharmacists
	Registration Inspection Units	Care in the Community, Homes (Group Homes), Nursing homes, Managers of independent sector establishments, Private hospitals, Clinics and hospices
Social Services Department	In-house services	Residential Care Homes (elderly, learning difficulties, mental health, physical disabilities, respite care), Day Centres, Home Care Services (in-house and purchased), Occupational Therapists, Children's Services, Special Schools, Other appropriate Local Authority departments (for example Education departments for equipment held in schools).
	Registration Inspection Unit	Any of the above services provided by the independent sector.

Table 14.3: MDA's Distribution Hierarchy for Safety Notices

The success of the MDA distribution hierarchy relies on individual Liaison Officers. They exercise discretion in disseminating particular warnings to appropriate managers and directors. The significance of the Liaison Officer is also acknowledged by the MDA in a range of practical guidelines that are intended to ensure the integrity of these distribution mechanisms. For instance, healthcare organisations must identify a fax number and e-mail address for the primary receipt of Hazard Notices and Device Alerts. They must also arrange for someone to deputise in the Liaison Officer's absence. The Liaison Officer is responsible for ensuring that Hazard Notices and Device Alerts are distributed immediately after publication. Safety Notices can take a less immediate route, as described in previous paragraphs. Liaison Officers are also responsible for documenting the actions that are taken following the receipt of Hazard Notices, Device Alerts and Safety Notices. In particular, they must record the recipients of these various forms of incident report. The documentation should also record when the reports were issued and a signed assurance from the recipient that any required actions have been taken.

Liaison Officers not only pass on Safety Notices to 'appropriate' managers, they can also choose to distribute particular warnings to staff. For instance, such direct actions might be used to ensure that new employees or contract staff are brought up to date with existing warnings. These groups of workers create particular problems for the distribution of incident reports in many different industries. Not only do they create the need for special procedures in the national system operated by the MDA, they also complicate the task of communicating recommendations from local systems. Changes in working procedures in individual hospital departments create significant training overheads for temporary 'agency' staff who may be transferred between different units over a relatively short period of time. When such training is not explicitly provided then it is likely that communications problems will occur during shift hand-overs [342].

Previous sections have argued that Liaison Officers play an important role within the particular distribution mechanisms that are promoted by the UK MDA. Aspects of their role are generic; they characterise issues that must be addressed by all reporting systems. For example, the conflict between the need for wide distribution and the problems of overloading busy staff apply in all contexts. Many reporting systems must also ensure that new workers and contract staff are brought up to date. Similarly, there is a generic tension between enumerating the intended recipients of a report and allowing local discretion to determine who receives a report. This last issue can be illustrated by the way in which particular, critical reports constrain or guide the actions of Liaison Officers. For example, a recent Device Bulletin into patient injury from bed rails explicitly stated that it should be distributed to all staff involved in the procurement, use, prescription and maintenance of bed rails. Liaison officers were specifically directed to ensure that copies of the report were forwarded to 'health and safety managers; loan store managers; MDA liaison officers (for onward distribution); nurses; occupational therapists; residential and nursing home managers; risk managers.' [538] In contrast, other Device Bulletins explicitly encourage Liaison Officers to adopt a far broader dissemination policy. A report into the (ab)use of single-use medical devices enumerated the intended recipients as all Chief executives and managers of organisations where medical devices are used, all professionals who use medical devices, all providers of medical devices and all staff who reprocess medical devices [536].

Previous paragraphs have focussed on the problems of ensuring that incident reports are disseminated effectively within the organisations that participate in a reporting scheme. We have not, however, considered the additional problems that arise when any lessons must be shared between organisations that operate their own independent reporting systems. Legislation is, typically, used to provide regulators with the authority necessary to ensure that safety-related information is shared through national or industry-wide systems. Such legal requirements often fail to address the concerns that many companies might have about providing information to such reporting systems. There is a clear concern that commercially sensitive information will be distributed to competitors. The exchange of safety-related information often raise questions about confidentiality and trust:

> "(The) FDA is keenly aware of and sensitive to the impacts of these new regulatory requirements on the pace of technological advancement and economic well-being of the medical device industry. At the same time, the agency is cognizant of the usefulness of information about the clinical performance of medical devices in fulfilling its public health mandate... FDA may require the submission of certain proprietary information because it is necessary to fully evaluate the adverse event. Proprietary information will be kept confidential in accordance with Sec. 803.9, which prohibits public disclosure of trade secret or confidential commercial information.." [252]

Less critical information, for instance about near-miss occurrences, may be retained within corporate reporting systems. Other organisations can then be prevented from deriving any insights that such reports might offer. The ability to overcome these barriers often depends upon the micro-economic characteristics of the particular industry. For instance, it can be difficult to encourage the altruistic sharing of incident reports in highly competitive industries. In other markets, especially those that are characterised by oligopolistic practices, it can be far easier to ensure the cooperation and participation of potential rivals. For example, the major train operating companies combined with the infrastructure provides to establish the CIRAS reporting system on Scottish railways [197]. This scheme has a lot in common with the CNORIS regional reporting system that has recently been established across Scottish NHS hospitals [417]. Another feature of these systems is, however, that the lessons are seldom disseminated beyond the small group of companies or organisations within the oligopoly.

The increasing impact of a global economy has raised a number of difficult moral issues that were not initially considered by the early proponents of reporting systems. For example, there have been situations in which the operators of a non-punitive reporting system have identified failures by individuals who work in counties that do operate punitive, legal approaches to adverse occurrences [423]. Such situations can create particular problems when individual employees may have contributed an incident report on the understanding that they were participating in a 'no blame' system. Although these dilemmas are relatively rare, it is important to acknowledge the increasing exchange of data between different reporting systems. For instance, the 49 MDA warnings, cited in previous paragraphs, resulted in 32 notifications being issued to other European Union member states [539].

The direct distribution of reports by the MDA to other EU member states represents one of several approached to the international dissemination of safety-related information. It effectively restricts the dissemination of information, in the first instance, to the other participants in the political and economic union. Other distribution mechanisms must be established on a country-by-country basis for the wider distribution of information, for example with the US FDA. The Global Aviation Information Network initiatives represent an alternative approach to the dissemination of safety-related incident reports [308]. As the name suggests, the intention is to more beyond regional distribution to provide global access to this safety information. Similar initiatives can be seen in the work of the International Maritime Organisation (IMO) and the International Atomic Energy Authority. Such distribution mechanisms face immense practical and organisational barriers. The same issues of trust and confidentiality that complicate the exchange of information between commercial organisations also affect these wider mechanisms. There is also an additional layer of political and economic self-interest when incidents may affect the viability and reputation of national industries. These problems partly explain the halting nature of many of these initiatives. They are addressing the *distribution problem* by making information available to many national and regional organisations. However, they often fail to address the *contribution problem* because very few reports are ever received from some nations.

14.2 From Manual to Electronic Dissemination

Previous paragraphs have argued that the problems of disseminating information about adverse occurrences and near miss incidents stem from the frequency and diverse range of publications; from tight publication deadlines and resource constraints and from the difficulty of ensuring that the intended readership can access a copy of the report. These problems have been addressed in a number of ways. For example, the last section examined a number of distribution models that are intended to ease the logistics of disseminating incident reports. The hierarchical approach adopted by the MDA was used to illustrate the manner in which key individuals, such as Liaison Officers, often lie at the heart of hierarchical approaches. In contrast, this section moves on from these organisation techniques to look at the way in which different technologies can be recruited to address some of the problems that complicate the dissemination of incident reports.

14.2.1 Anecdotes, Internet Rumours and Broadcast Media

It is important not to overlook the way in which information about an incident can be disseminated by word of mouth. This can have very unfortunate consequences. For instance, the U. S. Food and Drug Administration's Center for Food Safety and Applied Nutrition describe how the company at the centre of an investigation first became aware of a potential problem through the circulation of rumours about their involvement [253]. They report that 'the first news the dairy plant received that they were being investigated in relation to this outbreak was through rumour on the street'. The plant operators then demanded to know what was going on; 'Apparently someone had heard someone else talking about the Yersinia outbreak and how it was connected to the dairy plant'. These informal accounts then had to be confirmed with a consequent loss of confidence in the investigatory procedures that had prevented disclosure of the potential incidents before the rumour began.

Informal channels are often faster and, in some senses, more effective at disseminating information than more official channels. Rumours often circulate about the potential causes well before they are published by investigatory organisations. Very often official reports into an adverse occurrence or near-miss come as little surprise to many of the individuals who work in an industry. The dissemination of safety information by word of mouth is not entirely negative. Many organisations, such as the FDA and the MDA rely upon such informal measures given limited printing budgets and the vast audiences that they envisage for some warnings. Similarly, the use of anecdotes about previous failures has provided an important training tool well before formal incident reporting systems were ever envisaged or implemented.

There is a danger, however, that the information conveyed by these informal means will provide a partial or biased account of the information that is published by more official channels. Word of mouth accounts are likely to provide an incomplete view before the official report is distributed. This can also occur after the official publication of an incident report if individuals mis-understand or forget the main findings of an investigation. They may also be unconvinced by investigators' findings. In such circumstances, there is a tendency to develop alternative accounts that resolve uncertainties about the official report. These unauthorised reports are, typically, intended to gain the listeners' attention rather than to improve the safety of application processes. It is difficult to underestimate the impact of such informal accounts. They can undermine the listeners' confidence in the investigatory agency even though they may retain significant doubts over the veracity of the alternative account [278].

In recent years, the informal dissemination of incident related information has taken on a renewed importance. The growth of electronic communication media has provided significant opportunities for investigatory agencies to distribute 'authorised' accounts. The same techniques also enable engineers, focus groups and members of the general public to rapidly exchange information about adverse occurrences and near miss incidents. The recognition that e-mail, Internet chat rooms and bulletin boards can facilitate the 'unauthorised' dissemination of such information has attracted significant attention from organisations, such as the FDA. The issues surrounding these informal communications are extremely complicated. For example, there is a concern that drugs companies and device manufacturers might exploit these communication media to actively promote their products. This resulted in a recent initiative to directly consider the position of the FDA towards 'Internet Advertising and the Promotion of Medical Products' [256]. During this meeting, a representative of one pharmaceutical company argued that they had an obligation to make sure that the information available to the public was as accurate as possible. Given the lack of Internet moderation, however, it was impossible for companies to correct every misconception that might arise; 'we do not correct every piece of graffiti that may be painted in some remote area of Australia or Alabama or Philadelphia, but we do respond where we feel this is significant and we need to clarify the issues'. A representative of another drug company addressed rumours about adverse events more directly: "there may be a rumour that a certain product is going to be withdrawn at a certain time and if no one comes in and steps in who has a authoritative information and says, 'This is not true', that kind of rumour can absolutely snowball and can become uncontrollable if it is not quashed right when it starts" [256].

Companies are not the only organisations that can have a direct interest in refuting what can be termed *Internet rumours*. The FDA recently had to launch a sustained initiative to counter rumours about the safety of tampons [263]. The FDA identified three different versions of this rumour:

1. One Internet claim is that U.S. tampon manufacturers add asbestos to their products to promote excessive menstrual bleeding in order to sell more tampons. The FDA countered this rumour by stating that 'asbestos is not, and never has been, used to make tampon fibers, according to FDA, which reviews the design and materials for all tampons sold in the United States' [274].

2. Another rumour alleged that some tampons contain dioxin. The FDA reiterated that 'although past methods of chlorine bleaching of rayon's cellulose fibers could lead to tiny amounts of dioxin (amounts that posed no health risk to consumers), today, cellulose undergoes a chlorine-free bleaching process resulting in finished tampons that have no detectable level of dioxin'.

3. A final Internet rumour argued that rayon in tampons causes toxic shock syndrome (TSS) and could make a woman more susceptible to other infections and diseases. The FDA responded that 'while there is a relationship between tampon use and toxic shock syndrome–about half of TSS cases today are associated with tampon use–there is no evidence that rayon tampons create a higher risk than cotton tampons with similar absorbency'.

In order to counter these various rumours, the FDA launched a coordinated distribution of information on the Internet and to the broadcast media. This response indicates the seriousness with which they regard the Internet as a distribution medium for alternative or 'unofficial' accounts of particular incidents, in this case involving Toxic Shock Syndrome. Such actions do not, however, come without a price. They help to ensure that the public are aware of the scientific evidence in support of the FDA claims. They also inadvertently raise the profile of those Internet resources that disseminate the rumours in the first place. The FDA's response, therefore, adds a form of reflected legitimacy to the original arguments about the link between Tampon's and TSS. It is important to emphasise that our use of the term 'rumour' is not intended to be pejorative. In many cases, the informal dissemination of information can provide a useful corrective to the partial view put forward by more 'official' agencies. Such alternative sources of information must, however, support their claims and statements with appropriate warrants. In particular, it can be argued that these informal sources of information force official agencies to focus more directly on the issues and concerns that affect the general public. The Internet rumours about the relationship between tampons and TSS may have contained numerous statements that could not subsequently be supported, however, they did persuade the FDA to clarify the existing evidence on any potential links.

The previous case study illustrates some of the complex changes that are occurring in the manner in which information about adverse occurrences is being disseminated. Internet bulletin boards and chat rooms help to publicise rumours that are then picked up by the popular media. At this stage, regulatory authorities must often intervene to correct or balance these informal accounts. It is, however, insufficient simply to publish a response via an official web site which is unlikely to attract many of the potential readers who have an interest in a particular topic. The regulatory agency is, therefore, compelled to exploit more traditional forms of the broadcast media to refute rumours that were primarily disseminated via the web and related technologies.

This reactive use of the media represents a relatively recent innovation. More typically, investigatory agencies have used the press, radio and television in a more pro active manner to disseminate the findings of incident reports. As we have seen, this use of the media requires careful planning; there is a danger that the parties involved in an investigation may learn more about their involvement from the press than from more official channels. The FDA is similar to many national agencies in that it follows detailed guidelines on the use of the media to disseminate information. For example, media relations must be explicitly considered as part of the strategy documents that are prepared before each product recall. The dissemination of information in this manner must be treated extremely carefully. It is important that the seriousness of any recall is communicated to the public. It is also important to avoid any form of panic or any adverse reaction that might unduly influence the long term commercial success of the companies that may be involved in an incident. The sensitive nature of such recall notices is recognised in the FDA provision that the warnings may be released either by the FDA or by the recalling firm depending on the circumstances surrounding the incident [257]. The political sensitivity of these issues is also illustrated by the central role that is played by the FDA's Division of Federal-State Relations during Class I recalls. This classification is used when is expected and when the 'depth' of the recall is anticipated to require action by a retailers and consumers. The Federal-State Relations division is required to use e-mail to notify state and local officials of recalls that are associated with serious health hazards or where publicity is anticipated. These officials are then issued with enforcement papers that are prepared by the FDA Press Relations Staff. This mechanism illustrates the manner in which investigatory agencies may operate several parallel dissemination activities each with very different intentions. In addition to the publication of incident reports, press releases are prepared to initiate actions by the public and by retailers. These may be distributed at press conferences, by direct contact with particular reporters and by releases to all Associated Press and United Press International wire services. Further distribution mechanisms must also ensure that individuals within relevant organisations are 'well briefed' to respond to questions from the press.

It is also important to acknowledge the central role of press and media relations staff. Not only does this department warn other members of the organisation of media interest. They also ensure that their colleagues are adequately briefed to respond to media interest. Their ability to perform these tasks is dependent upon them being notified in the early stages of any incident investigation. FDA regulations require that the Press Relations Staff are notified by any unit that 'publicity has occurred relating to the emergency condition, as well as pending requests for information from the media and/or public' [257]. The senior media relations staff then liaise directly with the officials closest to the scene to ascertain what information needs to be released and when it should be disseminated to best effect. It can, however, be difficult to ensure that such press releases will be given the prominence that is necessary in order to attract the publics' attention to a potential hazard. Some warnings have a relatively high news value. The FDA's Consumer magazine often provides journalists with a valuable starting point for these incidents. For instance, a recent warning centred on a particular type of sweet or candy that had resulted in three children choking to death. Some of these products carried warning labels, suggesting that they should not be eaten by children or the elderly. Other labels warn of a choking hazard and say to chew the sweets thoroughly. Some were sold without any warning. This story attracted immediate and focussed media interest. Another warning, which was issued on the same day as the one described above, attracted far less media attention. This concerned the potential dangers of consuming a mislabeled poisonous plant called Autumn Monkshood [266] The packages containing the plant were mistakenly labeled with the statement, 'All parts of this plant are tasty in soup'. They should have indicated that consumption of the plant can lead to aconitine poisoning and that death could occur due to ventricular arrhythmias or direct paralysis of the heart. Simply releasing information to the media about potentially fatal incidents does not imply that all incidents will be equally news worthy nor that they will receive equal prominence in press, radio or television broadcasts.

As we have seen, it can be difficult for regulatory and investigatory agencies to use the media as a means of disseminating safety information. This involves the coordination of press releases and conferences. It also involves the training of key staff, such as press liaison officers, and the use of electronic communications techniques to ensure that other members of staff are informed how to respond to media questions. Even if this infrastructure is established there is no guarantee, without legal intervention, that a particular warning will receive the prominence that is necessary to attract public attention. Such problems are most often encountered by large-scale national systems. The issues that are raised by media dissemination of incident information are, typically, quite different for smaller scale systems. There can also be a strong contrast in media relations when incident information attracts 'adverse' publicity. This is best illustrated by the phenomenon known as 'doctor bashing' which has emerged in the aftermath of a number of incidents within the UK healthcare industries. Many professionals find themselves faced by calls from the government and from the media to be increasingly open about potential incidents. For example, Alan Milnburn the UK Health Secretary has argued that the "National Health Service needs to be more open when things go wrong so that it can learn to put them right" [111]. Together with this increased openness "they would also have to be accountable for their errors and prepared to take responsibility". Some doctors have described such statements and the associated media publicity as 'hysterical'. Recent BBC reports summed up this attitude by citing a General Practitioner from the North West of England; "Shame on the media for sensationalising and exaggerating incidents...shame on you for failing to report accurately adverse clinical events" [110].

Public and government pressure to increase the dissemination of information about medical incidents must overcome many doctor's fear of adverse or 'sensational' press coverage. At present, many NHS trusts have still to face up to the consequences of this apparent conflict. They are reluctant to disclose information about previous incidents even to their own staff for fear that details might 'leak' to the press. In this domain at least, we are a very long way from the culture of openness that the proponents of incident reporting systems envisage as a prerequisite for the effective implementation of their techniques. It is important not to simply view these tensions as simply the result of media interest in disseminating sensational accounts of adverse incidents. They reflect deeper trends in society. The chairman of the British Medical Association's Junior

Doctors' Committee saw this when he argued that "we have a more consumerist society... people are complaining more about everything... there is a lot of doctor-bashing in the press" [108]. Such quotations illustrate the way in which the media not only inform society, as in the case of FDA warnings, but they also reflect the concerns of society.

This section has focussed on the 'informal' dissemination of information about adverse incidents. In particular, it has focussed on the way in which electronic and Internet-based communications have provided new means of distributing alternative accounts of near-misses and adverse occurrences. We have also described how regulatory organisations have used the same means to rebutt these alternative reports. The conventional media is routinely used to support these initiatives. It can also be used to publicise more general safety warnings and can initiate investigations where other forms of reporting have failed to detect safety-related incidents. This more positive role must be balanced with the problems of media distortion that dissuade managers from disseminating the findings of incident reporting systems. There is a stark contrast between the use of the media to publicise necessary safety information and the fear of publicity in the aftermath of an adverse event.

14.2.2 Paper documents

The previous section has done little more that summarise the informal communication media that support the distribution of safety related information. Similarly, we have only touched upon the complex issues that stem from the role of the media in incident reporting. These related topics deserve books in their own right, however, brevity prevents a more sustained analysis in this volume. In contrast, the remainder of this chapter focuses on more 'official' means of disseminating incident reports. In particular, the following section analyses the strengths and weaknesses of conventional paper-based publications to disseminate safety-related information.

One of the most suasive reasons for supporting the paper-based dissemination of incident reports is to meet regulatory obligations. The importance of this media is clearly revealed in the various regulations that govern the relationship between the FDA, manufacturing companies and the end-users of healthcare products. The primary focus of these regulations is on the exchange of written or printed documentation. This emphasis is not the result of historical factors. It is not simply a default option that has been held over from previous versions of the regulations that were drafted in an age before electronic dissemination techniques became a practical alternative. As we have seen, the recommendations in some incident reports can have a legal force. Companies may be required to demonstrate that they have taken steps to meet particular requirements. This creates problems for the use of electronic media where it can be very difficult to determine the authenticity of particular documents. It would be relatively easy to alter many of the reports that are currently hosted on regulatory and governmental web-sites. Later sections will describe a range of techniques, such as the use of electronic watermarks, that can increase a reader's confidence about the authenticity of the documents that are obtained over the Internet. Unfortunately, none of the existing incident reporting sites have adopted this technology. In consequence, paper versions continue to exist as the 'gold standard' against which compliance is usually assessed. Copies obtained by other distribution mechanisms are, therefore, seens as in some way additional to this more traditional form of publication.

A further benefit of conventional, paper-based dissemination techniques is that regulatory agencies can exploit existing postal distribution services. A host of external companies can also be used to assist with the formatting, printing and mailing of these documents. The technology that is required to perform these tasks is well understood and is also liable to be readily available within most organisations. These are important considerations. Simplicity and familiarity help to reduce the likelihood of failures occurring in the distribution process, although as we have seen they are not absolute guarantees! Minimal staff training is needed before information can be disseminated in this way. It is for this reason that most small scale reporting systems initially exploit this approach. Typically, newsletters are duplicated using a photocopying machine and are then made available either in staff common areas or in a position that is close to a supply of reporting forms.

Paper-based dissemination techniques simplify the task of distributing incident reports because they can exploit existing mechanisms, including staff distribution lists as well as both internal and

Very Well	Well	Not Well	Not at All
17,862,477	7,310,301	4,826,958	1,845,243

Table 14.4: 1990 US Census Data for Self-Reported Ability in English

state postal services. There are further advantages. No additional technology, such as a PC with an Internet connection or CD-ROM, is required before people can access safety-related information. This is a critical requirement for the dissemination of some incident reports. One participant at a recent FDA technical meeting was extremely irritated by the continual reference to web sites as a primary communication medium. He asked the others present whether they knew how many American could access the Internet or could understand English [260]. Such comments act as an important reminder that paper-based publications continue to have an important role in spite of the proliferation of alternative dissemination techniques. For the record, Table 14.4 provides the latest available figures from the 1990 US Census describing self-reported English ability. The total US population was reported as 230,445,777 of which there were some 198,600,798 individuals who reported that they could only speak English. There were 31,844,979 who described themselves as being primarily non-English speakers. The self-reported figures for the standard of English amongst this community are shown in Table 14.4. The proportion of the population who express problems in understanding English appears to be relatively small. However, there may be a significant proportion of the population who did not return a census form and there is a concern that the proportion of non-English speakers might be relatively high in this community. There is also a natural tendency to over-estimate linguistic ability in such official instruments. Such factors motivate the provision of alternate language versions of safety-related information [822]. The 2000 census provided further insights into the growth of the Internet amongst the US population [823]. The census asked 'Is there a personal computer or laptop in this household?'. The returns indicated that 54,000,000, or 51%, of households had one or more computers in August 2000. This was an increase of 42% from December 1998 45,000,000, or 42%, of households had at least one member who used the Internet at hone, This had risen from only 26% in 1998 and 18% in 1997. Such statistics reinforce the point that significant proportions of the population in what is arguably the world's most technologically advanced nation still do not have Internet access. This is liable to be less significant for incident reports that are targeted at commercial organisations, for which one might expect a higher percentage of Internet connectivity. The census statistics are, however, a salient reminder for more general reports and warning such as those issued by the FDA that are deliberately intended for the general public.

Paper-based dissemination techniques are also resilient to hardware failures. It is a relatively simple matter to find alternative printing facilities and postal services. It can be far more complex to introduce alternative web-servers or automatic fax routing services. The reliability of the distribution service is only one aspect to this issue. There can also be considerable problems in ensuring that the intended recipients of incident reports can successfully retrieve alternative formats. Postal services are seldom swamped by the volume of mail. The same cannot be said by web servers or even by the use of fax-based distribution techniques. If the intended recipient's fax machine is busy at the time when an automated distribution service attempts to distribute an incident report, critical information can be delayed by hours and even days. At peak times of the day, many requests can either fail entirely or be significantly delayed as users request incident reports from the FDA or MDA web-sites. One particular problem here is that many government web sites only make limited use of more advanced techniques, such as predictive cacheing or mirror sites [416]. Similarly, the servers that provide access to incident reports may also be used to provide access to other documents that attract a large volume of users throughout the day. There is a certain irony in the manner in which some incident reporting web sites also elicit user-feedback about the failure of those sites that are intended to provide access to other forms of incident reports. Even if readers can download a computer-based report, there is no guarantee that they possess the application software that may be required to view it. Chapter 13 described how most incident reporting sites exploit either HTML and PDF. The former supports the dissemination of web-based documents because no additional support is required beyond a browser. Unfortunately, there is no guarantee that a document, which is formatted in HTML will be faithfully reconstructed when printed. This is significant because the psychological literature points to numerous cognitive and perceptual problems associated with the on-screen reading of long and complex documents [875]. In consequence, many organisations exploit Adobe's proprietary PDF format. PDF readers can be downloaded for most platforms without any charge. Problems arise, however, when incident reports that have been prepared for viewing under one version of the reader cannot then be viewed using other versions. For instance, a recent MDA report into Blood Pressure Measurement Devices contained the following warning: "Adobe Acrobat v.4 is required to view on screen the content of the tables at p.9 + 16...Adobe Acrobat v.3 can view remainder of document and can print in full". Paper-based dissemination techniques avoid such problems, which present a considerable barrier for many users who might otherwise want to access these documents.

There are further benefits to more traditional dissemination techniques. For instance, the physical nature of paper-based publications enables regulators to combine documents in a single mailshot. This is important because potential readers can skim these related items to see whether or not they are relevant to their particular tasks. This can be far more difficult to achieve from the hypertext labels that are, typically, used to encourage readers to access related items over the web [757]. The flexible nature of printed media can be illustrated by the way in which Incident Report Investigation Scheme news and safety alerts were directly inserted into printed copies of the Australian Therapeutic Goods Administration newsletter [45]. Similar techniques have been adopted by many different investigation schemes. Safety-related information is included into publications that are perceived to have a wider appeal. This is intended to ensure that more people will consider reading this information than if they had simply been sent a safety-related publication.

There are also situations in which investigatory and regulatory organisations have no alternative but to use printed warnings. For example, the FDA took steps to ensure that printed warnings were distributed about the danger of infection from vibrio vulnificus as a result of eating raw oysters [254]. The signs and symptoms of previous cases were described and the resulting warnings were posted at locations where the public might choose to buy or consume these products. The use of the Internet or of broadcast media provides less assurance that individuals who are about to consume raw Oysters are aware of the potential risks. This incident also illustrates some of the limitations of paper-based dissemination techniques. Many of the cases of infection were identified in and around Los Angeles. The FDA soon discovered that, as noted above, a significant proportion of this community could not speak or read English at the level which was required to understand the signs that had been posted. The States of California, Florida, and Louisiana only required Oyster vendors to post signs in English. In consequence, the FDA supplemented these printed warnings with a 24-hour consumer 'Seafood Hotline' that provided information in English and Spanish.

There are a number of problems that limit the utility of paper-based dissemination techniques as a means of distributing the documents that are generated by incident reporting systems. The most obvious of these issues is the cost associated with both the printing and shipping of what can often be large amounts of paper. These costs can be assessed in purely financial terms. They are also increasingly being measured in terms of their wider ecological impact, especially for large scale reporting systems that can document many thousands of contributions each year. Many organisations attempt to defray the expenses that are associated with the generation and distribution of incident reports by charging readers who want to obtain copies of these documents. This raises a number of complex, ethical issues. For example, the cost of obtaining a copy of an incident report can act as a powerful disincentive to the dissemination of safety-related information. This should not be underestimated for state healthcare services where any funds that are used to obtain such publications cannot then be spent on more direct forms of patient care. Some regulatory bodies, therefore, operate a tiered pricing policy. For example, the MDA do not make a charge for any of the Device Bulletins requested by members of the National Health Service. In contrast, Table 14.5 summarises the prices that must be paid to obtain copies of a number of recent MDA documents by those outside the national health system [540].

The costs illustrated in Table 14.5 do not simply reflect the overheads associated with the printing and shipping of these documents. They also, in part, reflect the costs of maintaining a catalogue of previous publications. This can prove to be particularly difficult with paper-based reports given the

Device Bulletins - 2001

Number	Title	Issue Date	Price
DB 2001(04)	Advice on the Safe Use of Bed Rails	July 2001	£15
DB 2001(03)	Guidance on the Safe Transportation of Wheelchairs	June 2001	£25
DB 2001(02)	MDA warning notices issued in 1995	May 2001	Free
DB 2001(01)	Adverse Incident Reports 2000	March 2001	Free

Device Bulletins - 2000

Number	Title	Issue Date	Price
DB 2000(05)	Guidance on the Purchase, Operation and Maintenance of Benchtop Steam Sterilisers	October 2000	£25
DB 2000(04)	Single-Use Medical Devices: Implications and Consequences of Reuse Replaces DB9501	August 2000	£15
DB 2000(03)	Blood Pressure Measurement Devices - Mercury and Non-Mercury	July 2000	£15
DB 2000(02)	Medical Devices and Equipment Management: Repair and Maintenance Provision	June 2000	£25
DB 2000(01)	Adverse Incident Reports 1999 Reviews adverse incidents reported during 1999 and describes MDA actions in response.	March 2000	Free

Table 14.5: Pricing Policy for Recent MDA Device Bulletins

storage that is required to hold the large numbers of publications that were described in the opening pages of this chapter. The MDA has published well over 300 different reports in the last five years. The logistics of supporting the paper-based distribution of such a catalogue has led many similar organisations to abandon such archival services. The Australian Institute of Health and Welfare now only provide the detailed back-up data and tables for many of their publications in electronic format [41].

A number of further limitations affect the use of paper-based dissemination techniques. The previous paragraphs have argued that such approaches do not suffer from the problems of server saturation and network loading that can affect electronic distribution mechanisms. Unfortunately, more convention dissemination mechanisms can suffer from other forms of delay that can be far worse than those experienced with Internet retrieval tools. Even with relatively efficient administration procedures there can be a significant delay between the printing of a report and the time of its arrival with the intended readership. These delays are exacerbated when safety managers or members of the general public require access to archived information about previous incidents. For instance, the MDA promise to dispatch requested reports by the next working day if they are in stock [541]. If they are not currently in print then they will contact the person or organisation making the request within forty-eight hours. These delays can be exacerbated by the use of the UK's second-class postal service to dispatch the requested copies of the report. This reduces postage costs, however, it also introduces additional delays. The second class service aims to deliver by the third working day from when it was posted. In the period from April to June 2001, 92.5% of second class 'impressions' satisfied this target. This is an important statistic because it implies that even if there is a relatively long delay before any requested report can be delivered, the duration of this delay is relatively predictable. In the same period, the UK postal servise achieved close to 100% reliability in terms of the number of items that were lost. The high volume of postal traffic does, however, mask the fact that Consignia received 223,495 complaints about lost items, 40,529 complaints about service delays and 37,256 complaints about mis-deliveries by the Royal Mail service from April to June 2001.

The delays introduced by a reliance on the postal service or similar distribution mechanisms also creates problems in updating incident reports. In consequence, organisations may be in the process of implementing initial recommendations at a time when these interim measures have already been revised in the final report. Updating problems affect a wide range of the publications that are produced from incident reporting systems. For instance, the FDA explicitly intended that their Talk Papers, which are prepared by the Press Office to help personnel respond to questions from the public, are subject to change 'as more information becomes available' [267]. Even when revisions are made over a longer time period, it is important not to underestimate the administrative burdens and the costs of ensuring that all interested parties receive new publications about adverse incidents. This point can be illustrated by the problems surrounding Temporomandibular Joints (TMJs). These implants have been used in several dental procedures. They were initially introduced onto the market before a 1976 amendment that required manufacturers to demonstrate that such products were both safe and effective. TMJs were, therefore, exempt from the terms of the amendment. From 1984 to June 1998, the FDA received 434 adverse event reports relating to these devices. 58% of these incidents resulted in injury to the patient. In 1993, the Dental Products Advisory Panel reclassified TMJs into their highest risk category (III). All manufacturers of TMJ devices were then required to submit a Premarket Approval Application, demonstrating safety and effectiveness, when called for by the FDA. In December 1998, the FDA called for PMAs from all manufacturers of TMJ implants. This was followed up by the publication in 1999 of a consumer handbook entitled, 'TMJ Implants - A Consumer Informational Update'. In April 2001 this was updated to present further information about the changing pattern of incidents involving these devices. As can be seen, adverse occurrences led to the publication of reclassification information in 1993. This had to be disseminated to all device manufacturers. This was revised in 1998 when the Premarket Approval Applications were called for. This change has considerable implications; the FDA have to ensure that they contact all of the commercial organisations that might be affected by such a change. TMJ's are relatively specialist devices and so only a hand-full of companies are involved in manufacturing them in the United States. It is important to recognise, however, that the Class III categorisation also applied to the sale of foreign imports. One solution to the potential problems that might arise in such circumstances is to use legal powers to require that all device manufacturers take measures to ensure that they are aware of any changes to the regulatory status of the devices that they produce. Such an approach is, however, infeasible for members of the general public and even for clinicians. It would clearly not be a productive use of FDA resources if their administrative staff had to answer repeated requests from concerned individuals who were simply wanting to check whether or not they had received the most recent information about particular devices.

The web offers considerable benefits for the dissemination of updated information about adverse occurrences and near-miss incidents. A single web-site can act as a clearing house for informations about particular products, such as TMJs, users can then access this page in order to see whether or not the information there had been updated. This approach raises interesting questions about the relationship between the reader and the regulatory or investigatory organisation that disseminates the information. In a conventional paper-based approach, a *push* model of distribution was used. The incident reporting organisation actively sent concerned individuals updated copies of information that they had registered an interest in. This enabled regulators to have a good idea about who read their reports. The overheads associated with this approach persuaded some organisations to adopt a *pull* model in which interested readers had to explicitly request particular documents. The dissemination of reports could then be targeted on those who actually wanted them rather than simply sending everyone a copy of every report. The administrative costs associated with such a scheme have persuaded many organisations to adopt the electronic variant of this approach in which individuals are expected to *pull* updated reports from a web page of information. This removes many of the costs associated with the production and distribution of paper-based reports. It also prevents regulators and investigators from having any clear idea of who has read the incident reports and associated publications that they have produced. Web server logs can prove to be misleading, given the prevalence of cacheing and other mechanisms for storing local copies of frequently accessed information [416].

14.2.3 Fax and Telephone Notification

Telephone and fax-based systems provide a compromise between the push-based approach of paper dissemination and the pull-based techniques of more recent, Internet approaches. In their simplest form, a pre-recorded message can be used to list all of the most recent updates and changes to paper-based documentation. This can help potential readers to identify the report that they want without consuming the regulator's finite administrative resources. It also enable frequent and rapid updates to be made to the information that is pre-recorded. Unfortunately, the linear nature of recorded speech can make this approach impractical for agencies that publish many different reports. A potential reader would have to listen to the recording for many minutes before hearing about a potential item of interest.

The use of pre-recorded messages to provide an index of updates to incident reports still does not address many of the administrative and resources problems that can arise from the paper-based distribution of these documents. At some point, copies of the report have to be printed and shipped to the prospective readers. One solution to these problems is to use fax machines to distribute incident reports. This approach has numerous benefits. For instance, fax-servers can be pre-programmed with large sets of telephone numbers. They will then automatically ensure that a faxed document is sent to every number of the list. More advanced systems will suspect a call if the fax machine is busy and will then re-try the number later in the run. The use of fax machines can also help regulatory authorities to keep track of the recipients of particular documents. For example, the UK MDA's institutional Business Plan for 2001-2002 includes the objective to monitor 'first time' fax failures when urgent safety warnings are issued to liaison officers. Of course, it is not possible to ensure that named individuals will have received and read a document that is sent in this manner. There can, however, be a reasonable degree of assurance that the fax has been received by the organisation associated with a particular fax number.

The FDA has pioneered the development and use of a more refined version of the systems described in previous paragraphs. They have developed a fully automated 'Facts on Demand' system. The user dials up the service and they then hear a series of instructions. If, for example, they press '2' on their keypad then they can hear more detailed instructions on how to use the system. If they press '1' then they can choose to order a document. If they dial 'INDX' or 4639 on the keypad then they can order an index of all documents on the system. If they choose to order an index, the system will call them back to fax a catalogue of publications. This currently runs to more than 50 page, however, it avoids the problems associated with listening to a pre-recorded listing for several hours! Callers can then use this faxed index to identify the identifier of the document that they want to retrieve. They must then call the system again, select the required option and then enter the document identifier. The system will then automatically fax them back with the required publication. The only technical requirement for the user of such a system is that they have access both to a fax machine and to a touch-tone telephone [264].

Most incident reporting systems continue to use paper-based dissemination techniques. Technological approachs, such as that described above, provide additional facilities that build on these more traditional approaches. This situation is gradually changing under increasing financial and administrative pressures. These influences can be seen behind the decision to move to the electronic publication of the FDA's User Facility Reporting Bulletin. In 1997, it was decided this it was no longer possible to print and mail this document out to anyone who requested it:

> "Time, technology, and budget restrictions have come together in the Food and Drug Administration. Ten years ago, our computer capability allowed us to communicate only within FDA. Now, with advanced computer technology we can globally communicate through the Internet and through Fax machines. As you would expect, Congressional budget cuts have affected all parts of government. FDA did not escape these cuts. In the search for ways to reduce our expenses, printing and mailing costs for distribution of publications in traditional paper form have come to be viewed as an extravagant expenditure... Now, budget restrictions prevent future distribution in paper form. We regret the need to move to this new technology if it means that many of our current readers will no longer have access to the Bulletin. We would like to remind you that you

> can also obtain copies through our Facts-on- Demand System or the World Wide Web." [867]

The concerns voiced in this quotation are understandable given the relatively low penetration of Internet connections into many areas of the US economy in 1997. Fax systems, such as 'Facts on Demand', provided an alternative dissemination technique. It can, however, be argued that they are likely to be replaced as more and more companies invest in Internet technology. Computers-based dissemination will then become the primary means of distributing incident reports. Before this can happen, however, we will need to address security concerns and the legal status of electronic reports. We will also need to consider the consequences of providing electronic access to large collections of safety-critical incident reports.

14.3 Computer-Based Dissemination

There are many diverse reasons that motivate the increasing use of information technology to support incident reporting systems. These approaches offer the potential for almost instantaneous updates to be disseminated across large distances. As we shall see, computer-based systems also offer security and access control facilities that cannot easily be provided using paper-based dissemination techniques. The same technological infrastructure that supports the rapid dissemination of individual incident reports also offers mass access to historical data about previous incidents and accidents. There are further motivations for providing this form of access to incident databases:

- *supporting risk assessment.* An important benefit of providing wider access to incident databases is that safety managers can review previous incidents to inform the introduction of new technologies or working practices. This information must be interpreted with care; contribution and reporting biases must be taken into account. Even so, incident databases have been widely used to support subjective estimates about the potential likelihood of future failures [423].

- *identifying trends.* Databases can be placed on-line so that investigators and safety managers can find out whether or not a particular incident forms part of a more complex pattern of failure. This does not simply rely upon identifying similar causes of adverse occurrences and near misses. Patterns may also be seen in the mitigating factors that prevent an incident developing into a more serious failure. This is important if, for example, safety managers and regulators were to take action to strengthen the defences against future accidents.

- *monitoring the system.* Regulators and safety managers can monitor incident data to determine whether particular targets are being achieved. Incident databases have been monitored to demonstrate reductions in particular types of incidents. They have also been used to support arguments about overall safety improvements. The following chapter will address the problems that affect this use of reporting systems. For instance, any fall in the number of contributions to a reporting system can reflect a lack of participation rather than an increased 'level' of safety.

- *encouraging participation.* If potential contributors can monitor previous contributions, they can be encouraged to participate in a reporting system. Information about previous incidents helps to indicate the types of events and near misses that fall within the scope of the system. Evidence of previous participation can also help to address concerns about retribution or of accusations about 'whistle blowing'.

- *transparency and the validation of safety initiatives.* Wider access to incident data helps to validate any actions that are taken in the aftermath of an adverse occurrences. For example, several reporting systems have used their incident data to draft a 'hit' list of the most serious safety problems [36]. By providing access to the underlying data that supports such initiatives, manufacturers and operators can see the justifications for subsequent regulatory intervention.

- *information sharing.* The development of on-line incident databases enables safety managers and regulators to see whether similar incidents have occurred in the past. This information technology provides further benefits. For the first time, it is becoming possible to extend the search to include the on-line databases of incidents in other countries and even in different industries. It is important not to underestimate the opportunities that this creates. For instance, it is possible to directly view incident data relating to the failure of medical devices in the USA prior to their approval for use in the UK. Conversely, it is becoming possible for the authorities in the USA to view elements of the submission for approval for particular forms of drugs that are submitted to the UK authorities. The electronic indexing of all of this data can help investigators, regulators and safety-managers to search through a mass of information that would otherwise overwhelm their finite resources.

The FDA recently summarised the advantages of electronic over paper based systems for incident reporting in the healthcare domain [271]. Firstly, they argued that automated databases enable readers to perform more advanced searches of information. This is important because it is likely that individuals may miss relevant information if they are expected to perform manual inspections of large paper-based data sets. Secondly, FDA also argued that computer-based retrieval systems can be used to view a single collection of information from a number of different perspectives. For instance, it is possible to present summaries of all incidents that relate to particular issues. This might be done by issuing a request to show every incident that involves a software bug or the failure of particular infusion devices. Similarly, other requests might be issued across the same collection of reports to identify incidents that occurred in particular geographical locations or over a specified time period. Such different views can, in principle, be derived from paper-based documentation of adverse occurrences and near-miss incidents. The costs of obtaining such information are, however, likely to be prohibitive. The third justification that the FDA identified for the use of electronic information systems was that they support the analysis of trends and patterns. Many incident reporting systems are investigation in new generations of 'data mining' applications and 'search engine' that can identify subtle correlations within a data-set over time. Finally, the FDA argued that electronic information systems avoid 'initial and subsequent document misfiling that may result from human error' [271]. As we shall see, this particular benefit can be more than offset by the problems that many users experience when they attempt to use computer-based systems to retrieve particular incident reports. Many applications require the use of arcane command languages or pre-programmed queries that are often poorly understood by the people that must use the information that is returned to them.

14.3.1 Infrastructure Issues

It is important not to automatically assume that all incident reporting systems are following a uniform path in the application of information technology. There is an enormous diversity of techniques. Some systems, such as the CIRAS rail application [197], deliberately avoid the use of computer networks. Security concerns have persuaded them to use stand-alone machines. This has profound concerns for the collation of distributed data. Paper forms are used throughout the Scots rail network and extensive use is made of telephone interviewing, even though it can often be difficult to arrange times when contributors can be contacted in this manner. Other organisations, such as the Swedish Air Traffic Control Organisation, have deliberately created computer-based reporting systems that exploit the benefits of networked applications. Individuals can log-onto the system from many diverse locations both to submit an incident report and to monitor the progress of any subsequent investigation. The following pages, therefore, review the strengths and weaknesses of the technological infrastructures that have supported incident reporting systems. The term 'technological infra-structure' is used here to refer to the means of distributing computer-based records of adverse occurrences and near-miss incidents. Subsequent sections examine issues that relate more to the retrieval of reports once they have been disseminated. In other words, this section looks at the techniques that investigatory bodies can use to push information out to end-users. Later sections look at the systems that end-users can exploit to search through that data and pull out information about particular incidents.

Stand-Alone Machines

Many incident reporting systems initially make very limited use of computer-based tools. Typically, they recruit mass-market desktop applications such as spreadsheets and text editors to help with managerial and logistical tasks. There are also more mature systems that have deliberately adopted the policy not to use more advanced computational tools. This decision can be justified in a number of ways:

- *security concerns.* The most pressing reason not to exploit more advanced technology in general, and network connectivity in particular, is that many safety managers have concerns about their ability to maintain the security of commercially sensitive incident data. Even if those operating the system have satisfied themselves that precautions can be taken against potential threats, senior levels of management may intervene to prevent the use of local or wide area networks. Security concerns are often most significant when an independent reporting agency holds data on behalf of operating companies. In such circumstances, the integrity and privacy of that information is often a prerequisite for running the system in the first place;

- *cost issues.* The declining costs associated with computer hardware have not been matched by similar reductions in connectivity charges within particular areas of the globe. In consequence, many small scale reporting systems may not be able to justify the additional expense associated with anything more advanced than a stand-alone machine. This is particularly important when the potential contributors to a reporting system are geographically distributed and may be involved in occupations that do not directly involve the use of information technology. For example, high wiring costs and legacy buildings have prevented many NHS trusts from providing direct network access to all of the wards in every hospital.

- *lack of relevant expertise.* Costs not only stem from the physical infrastructure. They also relate to the additional technical expertise that is required to connect stand-alone machines to local and wide area networks. It is relatively simple to register a single machine with an Internet service provider. Most incident reporting systems, however, depend on gather information from and disseminating information to large networks of contributors. Previous sections have also stressed the importance of maintaining connectivity within these networks to ensure that safety information is disseminated in a timely fashion. These factors combine to make it likely that any move beyond a stand-alone architecture will incur additional costs in terms of technical support to maintain the computer-based infrastructure.

- *'not invented here' syndrome.* The previous justifications for the continued use of stand-alone machines are well considered and appropriate. There is, however, a further reason for the longevity of this simple architecture that does little credit to the systems that embody it. As mentioned, many incident reporting systems begin by using commercial, off-the shelf packages to collate statistical data about previous incidents. As we shall see, many of the individuals who maintain the automated support are acutely aware of the problems and limitations that this software imposes on its users. There can, however, be a reluctance to move beyond these initial steps to elicit professional support from software engineers. This stems form a justifiable fear that the use of more advanced systems may imply a loss of control. In consequence, we see safety-critical data being held on mass-market systems whose licenses explicitly prohibit their use in such critical applications.

Stand-alone computers, without network connectivity, continue to play a significant role in many incident reporting systems. It is possible to identify a number of different modes of operation. They can be used simply to disseminate forms that can then be edited locally. Contributors can then print out the details of a particular incident and post the completed form back to a central agency. This is the approach advocated by the US Joint Commission on Accreditation of Healthcare Organisation's Sentinel Event system [431]. They request that all organisations transmit 'root cause analysis, action plan, and other sentinel event-related information to the Joint Commission through the mail' rather than by electronic means. The UK MDA also offer a range of electronic forms in PDF and in Microsoft Word format that can be printed and posted back when they have been completed [537].

The dissemination of incident reporting forms that are then intended to be printed and posted back to central agencies creates something of a paradox. Organisations that use stand-alone machines to coordinate their reporting policies must first obtain copies of these documents before they can complete their submission in a secure manner. The most common means of doing this is to use another networked machine to download the form then copy this across to the isolated machine on a disk. This is clearly a protracted mechanism. It may also fail to achieve the level of security that many people believe it ought to. For example, stand-alone machines are equally vulnerable to the free distribution of passwords and the problems associated with unlocked offices [1].

A slightly more complex use of information technology is to distribute a suite of programs that not only helps with drafting incident reports but also helps with the investigation and analysis of near misses and adverse events. Many of these systems are not primarily intended to support the exchange of information between institutions but to ensure that clear and coherent procedures are adopted within an institution. Rather than supporting the dissemination of information about incidents, regulators disseminate software support for local reporting systems. A recent collaboration between Harvard School of Public Health, the MEDSTAT Group, Mikalix & Company and the Center for Health Policy Studies illustrates this approach [3]. They devised the Computerised Needs-Oriented Quality Measurement Evaluation System (CONQUEST) system. This is intended to support general information about quality assurance in healthcare. It does, however, share much in common with many incident reporting systems. For instance, it was designed to help managers and clinicians derive answers to the following questions:

1. Did the clinician do the right thing at the right time?

2. Was effective care provided to each patient?

3. Was care provided safely and in an appropriate time frame for each patient?

4. Was the outcome as good as could be expected, given each patient's condition, personal characteristics, preferences, and the current state of medical science? [3]

The project began by devising a classification scheme for the information that it was to maintain. It then "became obvious that a computer database was the logical way to store and retrieve the data". A mass-market, single-user database application was then chosen as the implementation platform for this system. This decision reflects considerable expediency. It is possible to use these applications to quickly craft a working application that can be used to store information about several thousand incidents. This is sufficient for individual hospitals, however, it will not provide indefinite support as the number of incidents increases over time. Nor do these systems provide adequate support for the storage and retrieval of incident information on a national scale.

Further problems complicate the use of stand-alone architectures to support local incident reporting systems. Many of the distributed software systems that run on these applications offer means of tailoring the format of the data to meet local requirements. In general, this is an excellent approach because it provides safety managers with a means of monitoring incident-related information that might have particular importance within their working context but which has been ignored at a national level. Unfortunately, one consequence of this flexibility is that local systems often develop electronic data formats and classification schemes that are entirely inconsistent with those used by other organisations. These problems even occur when organisations exploit the same version of incident reporting software. Local changes can, over time, partition incident data so that it is difficult to join the individual approaches into a coherent overview of incidents across an industry. Technically, it is possible to match variant fields in different systems providing that it is possible to identify relationships between this information. This will not, however, resolve the problems of missing or partial data.

Electronic Mail

If incident information is only ever to be held and used locally then many of the previous caveats are of limited importance. In most cases, however, there is a need to support the exchange of incident information about near-misses and adverse occurrences. Arguably, the most common means of supporting such transmission is through the use of electronic mail. This generic term is usually applied to a range of software applications that exploit the Internet-based Simple Mail Transfer Protocol (SMTP). This enables users to send arbitrary messages between different accounts. The 'Seafood Network' provides a good example of the effective use of this simplest form of electronic communication [687]. This ensures that any message sent to central account, or list server, is automatically distributed to everyone who is registered with the service. The primary purpose of this "Internet based seafood network is to facilitate information exchange about the Hazard Analysis and Critical Control Point (HACCP) system of food safety control in the seafood industry". The HACCP system can be thought of as a form of incident reporting scheme that also disseminates more general safety-related information. There are, however, certain pitfalls that can arise from this *relatively* simple use of electronic dissemination techniques:

> "It has become apparent that some network users may not realise that the e-mail address, seafood@ucdavis.edu, automatically distributes a message to over 400+ subscribers worldwide. By all means, if you want everyone on the seafood network to read your message, address it to seafood@ucdavis.edu To communicate privately as a follow-up, please respond to the individual's e-mail address that is listed on the message.

As mentioned, SMTP supports the exchange of simple text messages. Some email applications also support Multipurpose Internet Mail Extensions (MIME). These enable senders to attach files of a particular type to their mail messages. This is useful if, for example, the sender of a mail message wanted to ensure that the recipient used CONQUEST or a similar system to open the file that they had attached to their email. MIME not only sends the file but also send information, that is usually hidden from the user, about those programs that can be used to open the attachment. It is clear that the recipients of any message must be able to interpret its contents. In the case of standard SMTP mail, the human reader must be able to understand the contents of any message. In the case of a MIME attachment, the associated program must be able to interpret the data that is associated with the mail message. Some industries are more advanced than others in specifying the format that such transmissions should take. In particular, there are many reasons why this issue should be of particular interest to healthcare professionals. There must be clear standards for the transmission of information if doctors are to interpret patient-related information sent from their colleagues. Similarly, the increasing integration of testing equipment into some hospital networks has created situations in which results can be automatically mailed to a clinician. In consequence, many professional bodies have begun to collaborate on standards for the transmission of clinical information. For example, Health Level 7 is an initiative to develop a Standard Generalised Markup Language similar to that used on the web, for healthcare documents [184]. These standards are not primarily intended to support incident reporting. They can, however, provide convenient templates for the transmission and dissemination of incident reports that are consistent with emerging practices in other areas of healthcare. This approach is also entirely consistent with attempts to develop causal taxonomies for incident reporting. The leaf nodes in techniques, such as PRISMA described in Chapter 11, might be introduced within such languages to record the results of particular investigations.

At present, most reporting systems do not exploit such general standards in the electronic transmission of incident information. Instead, they rely upon a range of formats that are tailored to particular reporting systems and its associated software. One of the most advanced examples of this approach is provided by the Australian Incident Monitoring System [36]. Initially, like many reporting systems, this relied upon the paper-based submission of incident forms from hospitals and other healthcare organisations. As the system grew, it established a network of representatives within 'healthcare units'. These representatives now collate the paper-based reports and enter them into a database that exploits the Structured Query Language (SQL). SQL can be viewed as a standard that describes the language that is used when forming requests for information. The software embodies the Generic Occurrence Classification (GOC). This supports the categorisation and subsequent analysis of the incident records that are held within the system. As mentioned, this initial data entry is performed within the 'health units'. At this stage, the system contains confidential information on those involved in the incident. It is, therefore, protected from legal discovery under Australian Commonwealth Quality Assurance legislation. For monitoring purposes, the Australian Patient Safety Foundation (APSF) then collates information from the individual units. Before this is done, all identifying information is removed from the individual reports. Current versions of the AIMS software enable individual units to email this information to the APSF system using the MIME techniques, described above. An important aspect of this transmission is that the individual records are encrypted prior to transmission. Later sections will describe how SMTP mail services are insecure and that such techniques are a necessary precaution against the unauthorised access to such information.

As mentioned, the AIMS approach is both innovative and well-engineered. There are, however, a number of potential problems with the techniques that it exploits. It adopts a model that is very similar to the stand-alone architecture, which was described in previous paragraphs. The collated database of anonymised incident information is held by the APSF as a central resources. This not only secures the data it also, potentially, acts as a bottleneck for other healthcare professional who might have a legitimate interest in analysing the data. In particular, safety managers in individual health units cannot directly pose queries to determine whether an incident forms part of a wider pattern. They must go through the mediation of the APSF. Increasing incident reporting systems are adopting a more egalitarian model in which anonymised incident data is also be distributed by electronic means. The move has been inspired by work in the aviation community and, in particular, the metaphor of an information warehouse that has been promoted by the GAIN initiative [308]. Those who contribute data should also have direct access to the data that is contributed by peer organisations. This egalitarian approach poses considerable logistical problems, including the difficult of ensuring the security of information transfers. As we shall see, a range of encryption techniques can be used during transfer. Password protection can also be used to restrict access. Neither of these techniques addresses the problems of ensuring the consistency of incident databases that may be replicated in each of the peer organisations. In other words, if the AIMS database were to be distributed more widely there would be a danger that one hospital might be using a collection that was updated in 2000 while others used more recent versions of the database. This problem arises because information about new incidents must not only be sent to the central clearing house operated by the APSF, it must also be sent to all peer organisations. Email can be used for this but there is no guarantee that every message will be acted upon and incorporated into the existing database. It is also likely that the size of many reporting systems would prevent any attempt to email out the entire collection at regular intervals. It is for this reason that organisations such as the National Transportation Safety Board (NTSB) exploit a mixture of on-line updates and CD-ROM digests of their incident databases. Organisations can then either download each new set of reports as they become available or simply order a CD-ROM of the entire updated collection.

CD-ROMS

Compact Disk-Read Only Memory provides a relatively cheap means of disseminating incident information without requiring that the recipient exposes their machine to the security risks associated with a network connection. This technology also avoids the level of technical support that can be associated with network administration. The popularity of is format is based on storage capacity of this media. Most CD-ROMs provide nearly 0.7 gigabytes of data; this is equivalent to almost 500 high-density floppy disks. The emerging successors to this format, such as Digital Versatile Disc (DVD-ROMs), expand this capacity to 8.5 gigabytes. CD-ROMs are also highly portable and light weight making the postal distribution of large amounts of data far cheaper using this format than printed documentation. It is also possible to encrypt the information on a CD-ROM; this provides added protection against the problems that can arise if critical documents go missing in the postal system. All of these technical attributes make this format particularly well suited as a communication medium for the dissemination of incident databases.

The CD-ROM format has further advantages. In particular, they provide a maximum data rate of between 2.8 and 6 megabytes per second on a 40x drive. This might seem a relatively trivial, technical statistic. However, such speed enable regulators and investigatory agencies to include multi-media resources, such as short audio and video clips, in addition to textual information and static images. Previous sections have described how CD-ROMs can be used to ensure consistency through periodic updates to the many different users of a reporting database. These databases do not, typically, make use of multimedia resources. In contrast, the additional facilities provided by the CD-ROM format tend to be exploited by some of the other publications that are generated by incident reporting systems. Hence, the sheer storage capacity of this medium provides means of disseminating textual incident databases. The access speeds supported by CD-ROM enable regulators to disseminate multimedia training presentations that are intended to guide safety managers and operators who must follow particular recommendations in the aftermath of near miss incidents and adverse occurrences.

The use of CD-ROMs to distribute information about adverse occurrences raises a number of further problems. In large, distributed organisations it may not be possible to provide all members of staff with access to personal computers that can play these disks. A number of innovative solutions have been devised to address this problem. One of these is illustrated by a staff training scheme that was developed by the Royal Adelaide Hospital. This scheme was closely tied to the Australian Incident Monitoring System, mentioned above. This hospital developed 'Mobi-ed' units that resemble the information booths that are found in public areas such as shopping malls and airports; 'the cabinet has solid wheels with brakes, a handle at the back for moving it around, and locks the computer behind two separate doors' [40]. They were based on standard desktop PCs. The booths also had the advantage that they could be left in common areas where a number of members of staff might have access to them during many different shift patterns. The units were deliberately moved between locations in the hospital; "this appearance and disappearance of the units encourages staff to check them out whenever they appear in their ward or work area" [40]. Multimedia training material was obtained to address specific training needs identified frm incident reports. The Mobi-Units also provided staff with access to an on-line tutorial about how and when to complete a submission to the AIMS reporting system.

It is important to recognise that CD-ROMs are simply a storage technology that supports the distribution of incident databases and training material. A number of deeper questions can, however, be raised about the effectiveness of the material that is distributed using this technology. Chapter 15 will analyse the effectiveness of incident databases. The following paragraphs provide a brief appraisal of multimedia training packages that are increasingly being developed by reporting agencies for dissemination on CD-ROM. We initially became interested in the effectiveness of this approach when developing materials for Strathclyde Regional Fire Brigade. Their training requirements are very similar to those of the healthcare professionals in the Royal Adelaide Hospital, mentioned above. Staff operate shift patterns and geographically distributed across a number of sites. There are also differences, for instance the training activities of a particular watch will be disrupted if they are called out in response to a call from a member of the public. Such circumstances increase the appeal of CD-ROM based training. The intention was to develop a series of multimedia courses that supported key tasks, which had caused particular concern during previous incidents. Fire-fighters could work through a course at their own pace. They could also suspend an activity and return to it after a call-out. There was also a perception that the use of computer-based technology might address the subjective concerns of staff who found conventional lectures to be ineffective. Figure 14.1 presents the results of a questionnaire that was issued to 27 staff within the Brigade, At the start of the project, they did not have access to any computer-based training. They received information about safety-related issues through paper publications, lectures and videos as well as drill-based instruction. However, many fire-fighters viewed 'real' incidents as an important means of acquiring new knowledge and reinforcing key skills. This finding has important health and safety implications. Incidents should reinforce training gained by other means. They are clearly not a satisfactory delivery mechanism for basic instruction.

As mentioned, a series of CD-ROM based training packages were developed to address the perceptions identified in Figure 14.1. These applications were developed in collaboration with the Brigade

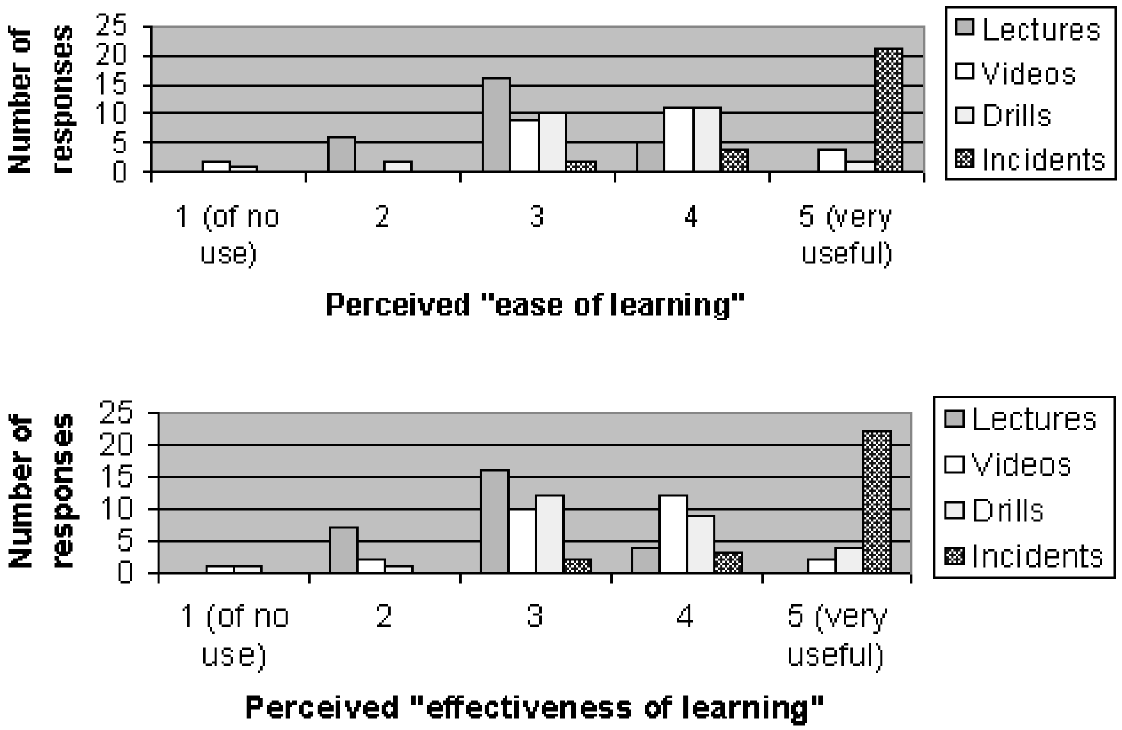

Figure 14.1: Perceived 'Ease of Learning' in a Regional Fire Brigade

training officer, Bill West, and two multimedia developers who were employed by the Brigade, Brian Mathers and Alan Thompson. Figure 14.2 illustrates one of these tools. This exploited the desktop virtual reality techniques, introduced in Chapter 8. Fire-fighters could use a mouse and keyboard to 'walk' into a Heavy Rescue Vehicle. They could then look inside equipment lockers and obtain a brief tutorial on the design and use of particular rescue devices. Previous incidents had illustrated the difficulty of providing officers with enough time to train on this particular vehicle given an operational requirement to keep it 'on call' as much as possible.

As mentioned, multimedia applications can be devised to address particular concerns that emerge in the aftermath of near miss incidents and adverse occurrences. The performance characteristics, in particular the access speeds of CD-ROMS, make this the favoured distribution media for such materials given network retrieval delays. It is important, however, to both understand and assess the various motivations that can persuade organisations to invest in tools such as that illustrated in Figure 14.2. In particular, we have argued that the strong motivational appeal of computer-based systems can support staff who find it difficult to be motivated by more traditional forms of training. We were, however, concerned that the introduction of computer-based techniques should not compromise particular learning objectives. We, therefore, conducted an evaluation that contrasted a computer-based system with more traditional techniques.

A matched subject design was adopted; each fire-fighter was paired with another officer of equivalent rank and each member of the pair was then randomly assigned to one of two groups. Both groups were given access to the same computer based training package on techniques to support the effective application of foam to combat particular types of fire. The technology used to produce the interactive application was similar to that used in Figure 14.2. One group was then given a CD-ROM based Comprehension Tool. This guided the officers through a series of questions about the training material and provided immediate feedback if any problems were diagnosed [425]. The

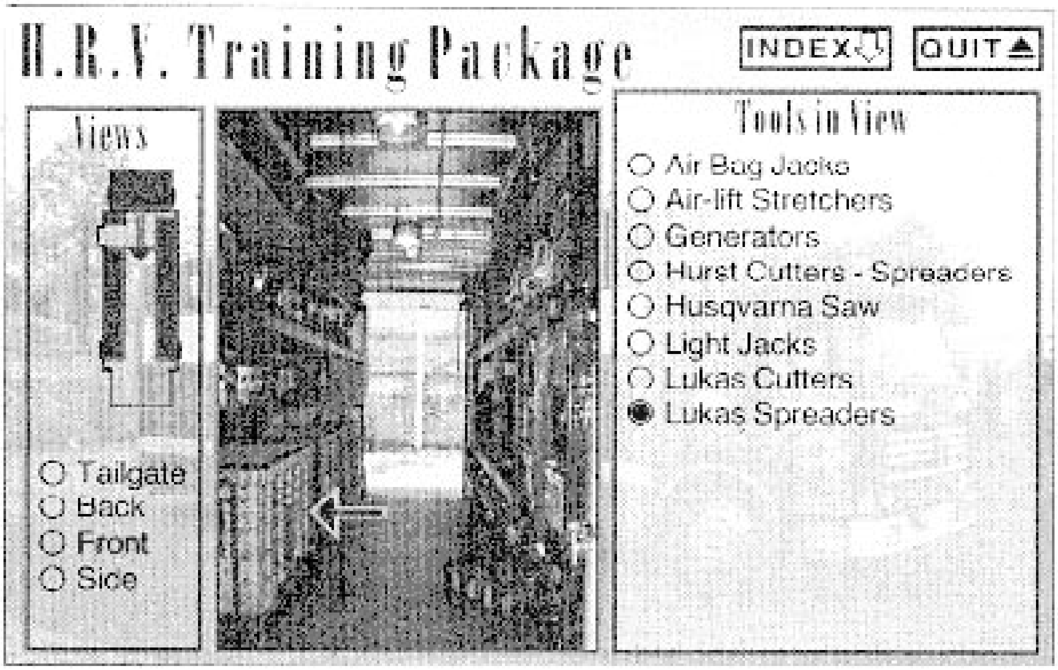

Figure 14.2: The Heavy Rescue Vehicle Training Package

second group was given a pencil and paper test without any feedback about the accuracy of their responses. One week later both groups were re-tested using the Comprehension Tool. It was hypothesised that the group that had previous access to the CD-ROM based self-assessment tool would achieve significantly higher scores than the group that had performed the pencil and paper test. A weakness in this experimental design is that learning effects might improve the results of the group that already had some experience with the Comprehension Tool. These effects were minimised by ensuring that both groups were entirely confident in the use of the tool before the second test began. Tables 14.6 and 14.7 present the results obtained for the two groups involved in this evaluation.

Rank	Number	Comprehension Tool or Paper test?	Score 1	Score 2
Sub officer	1	Paper	72	76
Leading fire-fighter	2	Comprehension tool	60	56
Leading fire-fighter	3	Comprehension tool	80	68
Fire-fighter	4	Paper	36	44
Fire-fighter	5	Paper	88	84
Fire-fighter	6	Comprehension tool	64	52
Fire-fighter	7	Paper	40	40
Fire-fighter	8	Comprehension tool	52	32

Table 14.6: Results for the first group of Fire Fighters

This evaluation forms part of a far wider attempt to validate the use of computer-based learning techniques. It does, however, provide a case study in the problems that can arise during these validation exercises. For instance, subtle differences in position 1 of the ranking schemes in Figures 14.6 and 14.7 complicate the task of making accurate comparisons. The station officer's performance is better than that of the sub-officer. There differences reflect the operating characteristics and compo-

Rank	Number	Comprehension Tool or Paper test?	Score 1	Score 2
Station officer	1	Comprehension tool	96	88*
Leading fire-fighter	2	Paper	80	60*
Leading fire-fighter	3	Paper	92	80
Fire-fighter	4	Comprehension tool	68	48
Fire-fighter	5	Comprehension tool	64	68
Fire-fighter	6	Paper	52	40*
Fire-fighter	7	Comprehension tool	84	64
Fire-fighter	8	Paper	60	60

Table 14.7: Results for the second group of Fire Fighters (* post fire)

sition of this organisation; they could not simply be changed for experimental expediency. Further problems arose because the second group of fire-fighters was called out while we administered the retest. It would have been unethical to prevent them from responding until after they had completed the evaluation! This study also provided some direct insights into the use of computer-based training techniques. Statistical T-tests failed to show any significant differences in the re-test scores between those who had access to the CD-ROM based tool and those who sat the pencil and paper test. We were unable to establish that the computer-based tool was any better than the more traditional techniques. Or put another way, it was no worse than existing methods for the dissemination of safety-related information.

The previous paragraphs are intended to correct the euphoria that often promotes the use of CD-ROM based technology for the publication of safety-related training materials. Often managerial and political pressures encourage the use of 'leading edge' technology without any careful analysis of whether this technology will support key learning objectives. Cost constraints can also act to limit many organisation's ability to disseminate the insights gained from previous incidents and accidents in any other format. Some regulatory and investigatory bodies have, however, continued to resist these pressures. For example, the State Training Teams of the FDA's Office of Regulatory Affairs support a 'lending library' of courses that must be presented by "trained state/federal facilitators that have already completed the original satellite course" [265]. These trained mentors are support by course videos, books, exams and answer key forms. It is interesting to note that these courses cover topics that are perceived to have a relatively high degree of importance in the FDA's regulatory role. For example, they include two courses on the investigation and reporting of adverse occurrences. This might imply that many of the issues addressed in previous chapters of this book cannot easily be taught using computer-based techniques.

Local and Wide Area Networks

The previous section has identified some of the strengths and weaknesses of CD-ROM technology for disseminating the information that can be derived from incident reporting schemes. They can be used to distribute the multimedia training resources that are intended to address previous failures. It can, however, be difficult to demonstrate the effectiveness of these resources. In contrast, CD-ROMs offer numerous benefits for the distribution of incident databases. They are relatively cheap. They offer relatively high storage capacity together with a relatively compact, lightweight format that is rugged enough to survive most postal services. Data can also be encrypted to provide additional security should a CD-ROM be lost or intercepted. There are, however, a number of limitations with this use of CD-ROM technology. In particular, it can be difficult to use this approach to issue more immediate updates to safety-related information. We have already describe the delays that can be introduced through the use of postal services to distribute physical media, such as CD-ROM. In contrast, many organisations are increasingly using computer networks to support the more rapid dissemination of information about adverse occurrences and near miss incidents. The MDA provide an example of this in their Business Plan for 2001-2002. They express the intention to develop closer links with the National Health Service and the Commission for Health Improvement with the objective of 'improving the dissemination' of information about adverse events. This will be done by 'electronic dissemination through our website and other Internet systems, so that healthcare professionals will increasingly have important safety information at their fingertips' [544].

It is convenient to identify two different sorts of incident data that can be accessed over computer networks. Firstly, incident databases help to collate information about large numbers of adverse occurrences and near misses. Secondly, incident libraries provide access to small numbers of analytical reports that may summarise the findings from many different incidents. In either case, these electronic documents must be stored in a particular format if they are to be disseminated across computer networks. Chapter 14 described the strengths and weaknesses of two of these formats. Hypertext Markup Language (HTML) documents can be viewed using standard browsers and are easily indexed by search engines but cannot easily be printed. Adobe's Portable Display Format (PDF) avoids this problem but most search engines have to be adapted to search this proprietary format for the keywords that are then used when users issue search requests. Incident data can also be stored in the file format that are supported by commercial mass-market databases and spreadsheets. This approach tends to be associated with incident databases. They are used to provide access to summary data about large numbers of individual incidents. PDF and HTML are more commonly used to support the dissemination of analytical surveys and the longer reports that are contained in on-line 'reading rooms'. The distinction between on-line libraries and incident databases is significant not simply because it influences file formats and retrieval techniques but also because it reflects important distinctions in the policies that determine what is, and what is not, made available over computer networks.

Private Databases and Public Libraries. Some organisations maintain private electronic databases that are not mounted on machines that are accessible to a wider audience. These same organisations may, however, provide wider access to the libraries of reports and recommendations that are derived from these private databases. This approach is currently being exploited by the MDA; 'we plan to introduce web reporting facilities that will feed directly into the Adverse Incident Tracking System (AITS) software' [544]. AITS is intended to help the Agency keep all its main records in electronic form for 'action and archiving'. It will also provide MDA staff with 'flexible data analysis tools to identify trends and clusters of incidents and that will enable us to adopt a more pro-active approach to reducing adverse incidents'. It is not intended that other organisations should have access to this database. In contrast, electronic access will be provided to what the previous paragraphs have characterised as 'libraries' of analytical overview documents. In passing, it should be stressed that the technical details of the AIMS software have not been released, not is there a detailed account of the full system development plan. It may very well be that the objectives outlined in the 2001-2002 Business Plan will be revised as AIMS is implemented.

Public Databases and Public Libraries. Other organisations provide access both to 'reading rooms' of analysis and to the databases of incidents that are used to inform these more analytical accounts. For example, the FDA's Manufacturer and User Facility Device Experience Database is freely available over the Internet [272]. It is comparable to AITS because it records voluntary reports of adverse events involving medical devices. An on-line search is available which allows you to search the Centre for Devices and Radiological Health's database of incident records. It is also possible to download the data in this collection over the Internet. These files are updated every three months. They are in a text format that enables safety managers and other potential readers to import them into a commercial database or word processor for further analysis. At the same time, the information in the MAUDE system is also used to inform more detailed incident investigations and surveys of common features across several adverse occurrences. The resulting reports are also available on-line in PDF format via an electronic 'reading room' [270]. This open dissemination policy enables readers to examine the warnings that are contained in a particular safety issue or alert publication. They can then trace additional details about particular incidents, and related occurrences, using the MAUDE database. It is important to stress, however, that the provision of public databases and reading rooms need not imply that sponsor organisations do not also maintain more private systems that are not made available in the manner described above.

Private Database and Private Summaries. Some incident reporting systems restrict access to both their database information and the reports that are derived from them. This policy is reflected in the way in which the AIMS system restricts access to its central database and only provides feedback on comparative performance to the individual units that participate in the scheme. These private summaries can be distributed over networks, either using the e-mail systems that have been described in previous paragraphs or using more explicit forms of file transfer [186]. This approach is entirely understandable given the sensitive nature of incident reporting within individual hospitals. There are other circumstances in which computer networks have been developed to support incident investigation and analysis within an organisation. The intention has never been that the data should be made public but that it should support specific tasks and objectives within the particular teams that must act upon incident data. This is illustrated by the U.S. Department of Health and Human Services' PulseNet system [259]. This system was established to distribute information generated from a molecular technique, pulsed-field gel electrophoresis (PFGE), that can be used to identify similarities between different samples of E. coli O157:H7. The PFGE technique was first used during a food-borne illness in 1993. This enabled laboratories in different locations to determine that they were fighting a common strain of bacteria. The lack of efficient computer networks to distribute the information from the independent PFGE tests prevented analysts from identifying these common features for the first week of the outbreak. Seven hundred people became ill and four children died in the outbreak. PulseNet is intended to reduce the interval taken to detect future incidents down to approximately 48 hours. This system is not intended to support the public dissemination of information about such events. It is, however, intended to support the analytical and decision making tasks that are necessary in order to detect common features between apparently isolated incidents.

It is important to emphasise that the increasing use of computer networks in incident reporting is only a small part of a wider move to intergrate diverse information sources about potential hazards. This integration is intended to support decision making. In other words, the dissemination of incident-related information is not an end in itself. This is illustrated by the manner in which epidemiologists can use PulseNet to trace common features in E. coli outbreaks. It is also illustrated by recent attempts to integrate diverse Federal databases to support the FDA field officers that have to determine whether or not to admit medical devices into the United States of America. The intention behind this initiative is to provide officers with rapid access to the range of data that they require in order to reach a decision. This data includes information about any previous adverse events involving particular products. However, this is only one part of a more complex set of requirements. For example, officers will also have to access the FDA's Operational and Administrative System for Import Support (OASIS) and Customs' Automated Commercial System (ACS). These data sources can be used to identify whether the product violates particular regulations by virtue of its point of origin. They can also be used to determine whether or not previous samples conformed with regulations when explicitly tested by the FDA. There is not intention that this integrated system should be widely accessible over public computer networks. It is, however, possible to access some of these information sources. For instance, it is possible to view the safety alerts that apply to particular products over the World Wide Web. This provides an example of the flexibility that such computer networks can offer for the provision of safety-related information. Users can choose whether to view warning that are sorted by particular industries, by country of origin or by FDA reference number. Users can also conduct free-text searches over the database of import alerts.

The FDA's import system contradicts the binary distinction between public and private distribution that was made in previous paragraphs. In practise, computer networks enable their users to make fine grained decisions about who can and who cannot access incident information. In the case of the import system, full functionality is reserved for field officers. Individual elements of the entire system are, however, made available for the public to access over the Internet. The same techniques can also be used more generally to restrict access to particular information about previous incidents. These approaches implement access control policies. The most common approach is to erect a 'firewall' that attempts to prevent access from anyone who is not within the local network that hosts the system. The following section discusses some of the consequences that such measures have for the implementation and maintenance of incident reporting systems. In particular, it is argued that compromises must often be made between restricting access and simplifying the procedures that users must follow in order to obtain access to incident data.

14.3.2 Access Control

Security deals with the unauthorised use and access to the hardware and software resources of a computer system. For example, *unauthorized disclosure* occurs when a individual or group can read information that they should not have access to. They, in turn, can then pass on information to other unauthorized parties. For instance, an unauthorised party might pass on information about an adverse occurrence to the press or broadcast media before that incident has been fully investigated. *Unauthorised modification* occurs when an unauthorised individual or group can alter information. They might have permission to 'read' data items but this does not automatically imply that they should also be able to modify data. It is, therefore, important to distinguish between different levels of permission. For example, an individual hospital contributing incident reports to a central database may have permission to access and modify their own reports. They might, in contrast, only be able to read reports from other hospitals without being able to modify them. Finally, *unauthorised denial of service* occurs when an individual or group can shut-down a system without authority for taking such an action. Unauthorised denial of service is a general problem in computer security. For example, the propagation of viruses can deny other applications of the computational resources that they require. I am unaware of any specific instances in which this form of attack has been a particular problem for incident reporting. It is important to stress, however, that unauthorised denial of service could have potentially profound consequences as incident reporting system become more tightly integrated into complex decision support systems, such as the FDA's Import application.

The issue of security affects incident reporting systems in a number of ways. For example, the Central Cardiac Audit Database (CCAD) project identifies two main threats [184]. The first centres on the security of data during transmission. When data is transmitted across open networks, such as the Internet, it can be intercepted unless it is encrypted. The second set of security concerns centres on controlling access to incident information after it has been collated. This is important because it is often necessary to ensure that different individuals and groups have different degrees of access to sensitive information. Some may be denied access to particular records. Other groups may be entitles to read data without being able to modify or 'write' it.

	Unit 1's Reports	Unit 2's Reports	Unit N's Reports
Administrator	read, write	read, write	read, write
Regulator	read	read	read
Unit 1	read, write	read	read
Unit 2	read	read, write	read
Unit N	read	read	read, write

Table 14.8: General Form for an Access Control Matrix

The distinction between 'read' and 'write' permissions has led to the development of access control policies. In their simplest form, these techniques implement a matrix that associates particular privileges with the users of a system and the objects that are held by that system. This is illustrated by Table 14.8. As can be seen, system administrators must be able to access and modify the reports that are submitted from all of the units that contribute to a reporting system. This requirement is, typically, enforced so that they can implement any revisions or updates that may subsequently prove to be necessary for the maintenance of the system. For example, they can automatically insert additional fields into the record of an incident. External regulators, in this instance, are provided with read-only access to all reports. Each of the contributing units can also read the reports from other contributors. They can also modify their own information. It is important to stress that the exact form of an access control matrix depends upon the nature of the reporting system. For example, some applications only provide read access to its contributors. They cannot modify their own data and all updates must be performed through an administrator who is entirely responsible for any 'write' actions. This reflects elements of the AIMS approach. In this case, the access control matrix would only contain write entries in the Administrator row. It is also important to emphasise that access is only granted if it is explicitly indicated in the matrix. By default, all other permissions are denied. In Table 14.8, the general public would not have any right to obtain or modify incident data.

Access control matrices are explicitly embodied within many of the more sophisticated software applications that have been developed to support incident reporting schemes. When a user makes a request to access a particular item of information, the system identifies the row associated with that user in the matrix. It then looks along the columns until it finds an entry associated with the object of the request. If the user does not have the relevant permissions then the request is denied. Unfortunately, this approach is not a feature of single-user systems. In general, access control makes little sense when there is only one row in the matrix. This has important consequences for many reporting systems that continue to use mass-market, desktop applications to support the dissemination of incident information. Single-user spreadsheets and databases, typically, have no means of making the fine grained distinctions implied by Table 14.8. In consequence, if a user is granted access to the system then they have complete permission to access all data. It is, typically, possible to apply locking techniques to the information that is held by these systems. This prevents unauthorised modification. However, this 'all or nothing' approach is usually too restrictive for large-scale systems [186].

Access control matrices provide numerous benefits to incident reporting systems. They explicitly represent the security policy that is to be enforced during the distribution of potential sensitive information. They are not, however, a panacea. As we have seen, it is entirely possible for individuals or groups to abuse their access permissions. For example, Unit 1 might pass on information about Unit 2 to a third party that is not entitled to this access, according to Table 14.8. In such circumstance, it is possible for system administrators to identify the potential sources of any 'leak' by inspecting the entries in the column that is associated with any disclosed information. Table 14.8 makes a number of simplifying assumptions. For instance, we have not considered 'grant privileges'. These enable particular users or groups to provide access permissions on certain objects. This is most often necessary when new Units join the system. The administrator would have to ensure that they were granted permission to read the contributions from all of the other Units. The entries associated with the system administrators would, therefore, be revised to *read, write, grant.* Paradoxically, the ability yo grant access also implies the ability to remove or deny access permissions. For instance, if incident information were being leaked to a third party then administrators might take the decision to remove all read permissions except those that apply to the Unit that contributed particular reports.

14.3.3 Security and Encryption

Access control matrices define the policy that is to be followed in the distribution and modification of incident information. In order to implement such a policy, most software systems rely upon encryption algorithms. As might be expected, these techniques take the original document, or plain text, and produce a cipher. Ideally, it should not be possible for an unauthorised person or group to derive the plain text from the cipher. One means of helping to prevent this is to create an encryption algorithm that relies not only on the plain text but also an additional input parameter known as a key. This can be illustrates by Caesar's algorithm. Caesar's algorithm replaces each letter in the plain text with the next letter in the alphabet. The letter 'a' would be replaced by 'b' in the cipher, 'c' would be replaced by 'd' and so on. This is a relatively simple algorithm to guess and so we might require that the user also supplies a key. The key could be the number of places that each letter is offset. For example, in order to decipher the following phrase we must know that each letter has been shifted by 14 places in the alphabet:

Cipher:	VAPVQRAG ERCBEGVAT
Plain text:	INCIDENT REPORTING

This simple algorithm is vulnerable to many different forms of attack. For instance, we can examine letter frequencies in the cipher to make guesses about the identity of particular letters. It does, however, illustrate the key features of *secret or private key* encryption. In order for this approach to work, it is important that the key is never disclosed to unauthorised individuals. The previous example also illustrates further aspects of secret key encryption. For instance, this approach can be strengthened by choosing a suitably complex algorithm. Alternatively, it is often more convenient to choose a relatively simple algorithm but a very complex key. Without the key, even if the algorithm is well known, it can be extremely difficult for unauthorised individuals to decipher a message. It is for this reason that encryption software often emphasizes the size of the keys that they support, 64 or 128 bits for example.

Secret key encryption is a feature of many national reporting systems. Regulatory and investigatory agencies provide each unit in the system with the encryption software and a password. This is then used to protect incident reports when they are transmitted across computer networks. This approach was adopted by the APSF's AIMS system when they moved from paper to electronic submissions [35]. One problem with secret key encryption is that it can be difficult to secure the dissemination of keys. They cannot be transmitted over the computer network because this would defeat attempts to secure transmission over that network. Conversely, it is impossible to secure the network until the secret key has been agreed upon.

Public key encryption provides an alternative to secret key encryption. This relies upon algorithms that require different keys to encode and decode the plain text. Typically, users distribute a public key to anyone who might want to send them secure information. They can do this because they know that this key can only be used to encode data. Only they have access to the second private key that is required in order to read any message. This is the approach adopted by the 'Pretty Good Privacy' or PGP mechanisms that are widely available over the Internet. PGP is one of several systems that are recommended for transmission of data to the Central Cardiac Audit Database, mentioned earlier [184]. The PGP package provides a variety of utilities for the generation, management and use of encryption keys. It has the advantage of being low or no cost and is widely available. Secure/Multipurpose Internet Mail Extensions (S/MIME) is an application of public key encryption to the MIME technology that is widely used for the dissemination of incident data. This approach is supported by recent versions of Netscape Communicator and Internet Explorer. S/MIME applies encryption to individual files that are mailed from one machine to another. At is also possible to create secure links that encrypt information passing between two or more machines. This approach is deliberately designed to support more interactive forms of communication, such as web browsing, where information can be passing in both directions over the connection. The best known application of this approach is known as the Secure Socket Layer (SSL) protocol. This applies public key encryption over an entire session rather an individual item of mail.

PGP, S/MIME and SSL have all been used to secure the data that is transmitted between the contributors of incident reports and regulatory or investigatory agencies [271, 184]. It is important to emphasise that these technologies do not provide any absolute guarantees about the security of any electronic communication. It is theoretically possible to break most implementations. The technical expertise and computation resources do make it extremely unlikely that this will occur, at least in the short term. These observations re-iterate an important concept; it is seldom possible to achieve absolute security. Safety managers must adopt an informed approach to risk assessment. The degree of technological sophistication applied to secure incident data must be proportional to the sensitivity of that data. It is important, however, that these issues are explicitly considered as more and more reporting systems use computer-based networks as a cheap and effective means of disseminating information about near misses and adverse occurrences [378].

One way in which cryptography has been used by incident reporting systems is to support digital signatures. A digital signature is a means of encoding a message in such a manner that it authenticates the sender's identity. This is important if regulators and investigation agencies are to ensure that reports of an incident have not been sent for malicious reasons. Both private and secret key techniques can be used to implement digital signatures. The fact that the recipient can decode a message that was encoded using the secret key agreed with the sender might, at first sight, seem to be sufficient for a secret key implementation. No other person should know the secret key. Unfortunately, this is vulnerable for a number of reasons. This approach is vulnerable to a replay attack in which a previous message is saved and later resent by some unauthorised agent. Similarly, some portion of a previous message may be cut and pasted to form a new message. This is feasible because it is possible to make inferences about the contents of a message even even if it is impossible to completely decipher all of its contents. For these reasons, secret key implementations of digital signatures usually also encode characteristics of the entire message, such as the date when it was sent and the number of characterise in the plain text. When the message has been decoded the recipient can check this additional information to ensure the integrity of the content. This technique can also be used when the message is not, itself, encoded. A signature block can be encrypted at the end of the message. Again, the recipient can decode the signature block and use the techniques described above to establish its authenticity. This approach has given rise to a range of more elaborate techniques that support the concept of *electronic watermarks* .

Public key implementations of a digital signature can be slightly more complicated. For instance, a contributor might encrypt a message using it's secret key. It will then encrypt the results of this encryption using the public key of the regulator. The message is then sent over a computer-based network. The regulator first decodes the message using their secret key. Ideally, no other users can complete this first step assuming that the regulators secret key is not compromised. Next, the regulator can apply the contributor's public key to extract the plain text. The regulator knows that the contributor sent the message because only the contributor has access to their secret key.

It might seem that such details are a long way removed from the practical issues that must be considered in the development and operation of incident reporting systems. The key point about digital signatures is that they enable organisations to transmit information in a secure manner that can be granted the same legal status as conventional, paper-based documents. For instance, in 1997 the FDA issued regulations that identified the criteria that would have to be met for the use 'of electronic records, electronic signatures and handwritten signatures executed to electronic records as equivalent to paper records and handwritten signatures executed on paper' [271]. These regulations applied to all FDA program areas and were intended to support the widest possible use of electronic technology 'compatible with FDA's responsibility to promote and protect public health'. The FDA requirements illustrate the importance of understanding some of the concepts that have been introduced in previous paragraphs. For instance, Section 11.70 requires that 'electronic signatures and handwritten signatures executed to electronic records must be linked to their respective records so that signatures cannot be excised, copied, or otherwise transferred to falsify an electronic record by ordinary means'. Such concerns have motivated the development of message-specific signature blocks mentioned above.

There are many reasons why incident reporting systems are forced to introduce some of these more advanced security measures. The FDA regulations reflect a concern to ensure that electronic reports of adverse occurrences have the same legal status as their paper-based counterparts. These more advanced security techniques can also be implemented in response to the concerns expressed by potential contributors. There is often fear that an individual's identity will be revealed. Similarly, safety managers in the healthcare industry must also respect patient confidentiality. There is, however, usually a trade-off to be made between the security of a system and the ease with which it can be used by its operators. For example, the use of cryptographic techniques implies a considerable managerial overhead both on the part of systems administrators and the users who must remember the passwords that protect and sign their information. Recent studies have suggested that for every user of a system there is a request for support staff to reset a forgotten password every three-four months [1].

It is difficult to under-emphasise the human element in any secure system. For example, recent attempts to introduce public and private key cryptography into one national reporting system produced a series of responses from potential contributors. They commented that these human issues would compromise the most advanced technology. It is possible for an owner to 'lend' a file, as a collaborative fraudulent gesture, or to unwittingly assist a fraudulent colleague in an 'emergency'. The FDA has acknowledged that 'such fraudulent activity is possible and that people determined to falsify records may find a means to do so despite whatever technology or preventive measures are in place' [271]. There are also more 'mundane' threats to the security of incident reporting systems. Previous research has suggested that people are often very lax in their selection of passwords [880]. The use of recognisable words, of names of friends, of addresses makes the entire system vulnerable to dictionary attacks. These take the form of repeated requests to access a system where a different word or phrase is supplied from an electronic dictionary in response to each password request. Eventually this form of attack will succeed unless care has been taken in selecting a password. Further problems arise when passwords are distributed to friends and colleagues. More commonly, however, systems are compromised by simply leaving a machine connected to the network. Users issue the requested password and then leave the machine unattended. I recently saw an example of this on a hospital visit. A screen-saver was used with the warning 'unauthorised access on this machine is prohibited'. Anyone who ignored this message could have directly accessed and altered the patient records for that ward while the nurse was away from her station.

Mnny of the usability problems that affect secure systems can be reduced through the use of biometric authentication. These techniques avoid the need for users to remember arbitrary passwords. They include the use of fingerprints, retinal patterns, signatures and voice recognition [441] This technology has not yet been widely used to secure the transmission of incident data. It is, however, likely that it will be used within future systems. Part of the reason for this is apparent in recent observations made by the Central Cardiac Audit Database Project. They argue that such security techniques appear to be 'an extremely expensive solution to a non-problem, but public fears about Internet security, mostly unfounded but encouraged by the popular media, would need to be allayed before widespread medical data transmission via the Internet would be acceptable' [184].

14.3.4 Accessibility

'Accessibility' can be thought of as the converse of access control. Just as it is important to ensure that unauthorised people are denied access to an incident report, it is equally important that authorised individuals can obtain necessary information. This implies that any computer-based resource should be evaluated to ensure that the human-computer interface does not embody inappropriate assumptions about the potential users of such as system. This is particularly important because the computer-based dissemination of incident reports can offer particular advantages to certain groups of users providing that their requirements are considered during the early stages of systems development. People with visual disabilities can use a range of computer-based systems to access incident databases in a manner that cannot easily be supported using paper-based techniques. Unfortunately, the use of icons and complex menu structures can prevent many users from exploiting screen readers and similar devices. It is for this reason that many organisations publish minimum standards in this area, such as Federal Regulations Section 508 on the accessibility of electronic information. In order to satisfy these requirements it is important that the developers of incident reporting systems provide some means for users to communicate any difficulties they might have experiences in access their data. This can, however, create a recursive problem in which the users of the information resource cannot access the information resource in order to learn of alternative format or other forms of help:

> "The U.S. Department of Health and Human Services is committed to making its web sites accessible to all users. If you use assistive technology (such as a Braille reader, a screen reader, TTY, etc.) and the format of any material on our web sites interfere with your ability to access the information, please use the following points of contact for assistance. To enable us to respond in a manner most helpful to you, please indicate the nature of your accessibility problem, the preferred format in which to receive the material, the web address of the requested material, and your contact information..." [31]

Although Section 508 of the US Accessibility Act focuses on users with special needs, there is also a more general requirement to ensure that people can access incident information. In consequence, observational studies and laboratory-based evaluations may be conducted to ensure that users can operate computer-based information systems. This implies that designers must consider the previous expertise of their users and of their ability to exploit particular human-computer interaction techniques. Brevity prevents a more detailed introduction to the design and evaluation of interactive computer systems in general. Preece et al provide a survey of techniques in this area [686]. In contrast, the following pages focus more narrowly on techniques that can be used to identify patterns of failure in large-scale collections of incident reports.

14.4 Computer-Based Search and Retrieval

Previous sections have considered the dissemination of information about individual incidents. In contrast, this section focuses more narrowly on the problems that arise when providing access to databases of previous incidents over computer networks. Recent technological innovations, often associated with mass-market applications of the World Wide Web, are creating new opportunities for rapidly searching large number of incident reports that can be held in many different countries across the globe. Before looking in more detail at these 'leading edge' systems, it is first necessary t understand why organisations are exploiting computer-based dissemination techniques for their incident databases.

The sheer scale of many reporting systems motivates the use of electronic dissemination techniques for incident databases. The number of incident reports that are submitted to a system can accumulate rapidly over a relatively short period of time, even in local or highly specialised systems. For instance, the Hyperbaric Incident Monitoring Study was started in 1992 to collects reports of incidents and near misses in hyperbaric medicine. Hyperbaric Oxygen Therapy is the most common form of this treatment. The patient enters a chamber that is filled with compressed air until a required pressure is reached. The patient breaths 'pure' oxygen through a mask or a transparent hood. In addition to diving recompression, this techniques has been used in the treatment carbon monoxide poisoning, wound healing and post radiation problems. This system is currently operated through the APSF, the same organisation that maintains AIMS. The Hyperbaric Incident Monitoring Study was launched internationally in 1996 and the associated forms have been translated into 4 different languages. By early 2001, there were some 900 reported incidents in the database [38]. This partly reflects the success of this system. It also creates considerable practical problems. The costs associated with maintaining even a relatively simple paper-based indexing system would be prohibitive. It would also be difficult for other organisations to access this data without replicating each paper record or posing a succession of questions to the staff who are responsible for maintaining the paper indexing system.

It is feasible but unlikely that the Hyperbaric Incident Monitoring Study database could be implemented using paper-based techniques. In contrast, other national reporting databases could not be maintained without electronic support. Firstly, the sheer volume of reports makes it essential that some form of database be used to collate and search that data. Secondly, the large number of individuals and groups who might legitimately want to retrieve incident information increase the motivation to provide access to these databases over computer networks. For instance, the FDA's Adverse Event Reporting System (AERS) for adverse events involving drugs and therapeutic biological products contains more than 2 million reports. Table 14.9 provides a break-down of the number of incidents that are entered into the FDA's databases within a single year. The Center for Biologics Evaluation and Research's Error and Accidents Reporting System records incidents that occur in the manufacture of biological products. An error or accident is a deviation from the 'good manufacturing practice' set down by FDA regulations. The Drug Quality Reporting System receives reports of similar incidents that affect the manufacturing or shipping of prescription and over-the-counter drug products. These incidents can result in problems for the formulation, packaging, or labeling of these products. Post-marketing surveillance for vaccines is handled by the Vaccine Adverse Event Reporting System. Approximately 15 percent of the reports describe a serious event, defined as either fatal, life-threatening, or resulting in hospitalization or permanent disability. The Manufacturer and User Device Experience (MAUDE) Database receives between 80,000 and 85,000 reports per year. The 1984 Medical Devices Reporting regulation required manufacturers to report device-related adverse events to the FDA. In 1990, the Safe Medical Devices Act extended this regulatory structure to include user facilities such as hospitals and nursing homes. Serious injuries that are device-related must be reported to their manufacturers. Fatalities must be reported both to the manufacturer and directly to the FDA. The MAUDE database was established in 1995 to support the Safe Medical Device Act and now contains more than 300,000 reports. Another 500,000 reports are collected in a pre-1995 database. Finally, table 14.9 records that the FDA's risk-based summary reporting system receives some 30,000 reports per annum. Products that are approved for this summary reporting process are 'well known' and have a 'well-documented' adverse event history [276]. This approach involves the periodic submission of adverse event statistics in a tabular form that yields 'economies for both the devices industry and FDA'.

Reporting System	No. of reports per annum
Adverse Event Reporting System	230,000
(CBER) Biologics Error and Accidents Reporting System	13,000
Drug Quality Reporting System	2,500
Vaccine Adverse Event Reporting System	12,000
Manufacturer and User Device Experience	80,000
Risk-Based Summary reports	30,000

Table 14.9: Annual Number of Incidents Included in FDA Databases

The volume of data that can be gathered by successful national systems justifies the use of information technology. However, the use of this technology does not provide a panacea for the management of large-scale databases. For instance, the cost implications and the requirement to use specialist hardware and software often convinces many public or Federal agencies to involve contractors to run these systems. The FDA's Drug Quality Reporting System database, mentioned above, is run in this manner. FDA staff interact with the system via an on-line interface that is intended to help them pose particular queries or questions that are then relayed to the database, which is administered by the contract organisation. The management of the Vaccine Adverse Event Reporting System database is even more complex. This is jointly administered by the FDA's Center for Biologics' Division of Biostatistics and Epidemiology and the Centers for Disease Control and Prevention, Vaccine Safety Activity, National Immunization Program. Representatives of both agencies oversee data processing and database management that is again performed by a contractor. Such complex relationships occur in other industries. For instance, the ASRS is largely funded by the FAA. NASA manages the system, which is in turn operated under contract by the Battelle Memorial Institute. In most cases, these relationships have proven to be highly successful. There are, however, considerable problems in ensuring that technological requirements are accurately communicated between all of the stake-holders in such complex, interactive applications. In consequence, previous studies have revealed considerable frustration from users who feel that many incident databases no longer support all of the retrieval tasks that they must perform [472, 413].

Further problems affect the management of large-scale incident databases. For example, it is often convenient for national regulators to establish a number of different schemes that focus upon particular incidents. For instance, Table 14.9 summarises the different schemes that are operated by the FDA, This can create problems because the same incident can fall within the scope of more than one database. This problem can be exacerbated if different agencies also run apparently complementary reporting systems. For instance, the FDA has to monitor medication error reports that are forwarded by clinical staff to the United States Pharmacopeia and to the Institute for Safe Medication Practices. Some of these incidents are then incorporated into the FDA's own databases. Similarly, the FDA must also review the medical device reports that are submitted to the MEDWATCH programme in case they have any relevance for possible medication errors. In addition to all of the systems mentioned in Table 14.9, the FDA also maintains a central database for all reports involving a medication error or potential medication error. This contains some 7,000 reports. In total contrast to this amalgamated database, the FDA's Vaccine Adverse Event Reporting database is entirely 'independent of other FDA spontaneous reporting systems' [276]. This diversity creates a flexible approach that is tailored to the various industries which are served by the FDA. It also creates considerable managerial and technical problems for those individuals who must support the exchange of data both within and between the various systems. For instance, some of these systems provide public access to incident data. There is a web-based search engine that can be used to find the detailed records held in the MAUDE system. Other applications are strictly confidential and no public access is provided. This can create problems if incident data is transferred from one system to the other. Confidential reports can be made public if they are transferred to the open system. Alternatively, if potentially relevant reports are not disclosed then valuable device-related safety information will be withheld and the MAUDE data will be incomplete. This would create doubts about the value of the system as a means of tracing more general patterns from information about previous incidents.

Further legal and ethical issues complicate the use of on-line systems to help search through incident databases. For example, a number of countries now have powerful disclosure laws that enable people to access databases in the aftermath of near miss incidents and adverse occurrences. This provides an opportunity to search through the database and identify previous failures with similar causes. In subsequent litigation, it might then be argued that responsible organisations had not shown due care because they had failed to learn from those previous incidents. Several safety managers have described their concerns over this scenario during the preparation for this book. One even commented that their company was considering deleting all records about previous incidents. These concerns are partly motivated by the difficulties that many organisations have in storing and retrieving information about previous incidents. Incidents often recur not because individuals and organisations are unwilling to learn from previous failures but because they lack the necessary technological support to identify patterns amongst thousands of previous incident reports. For instance, many users of incident databases cannot accurately interpret the information that is provided by database systems. It is also difficult for safety managers to form the commands that are necessary to retrieve information about particular types of incident. The following sections describe these problems in greater detail and a number of technological solutions are proposed.

14.4.1 Relational Data Bases

There are two central tasks that users wish to perform with large-scale incident databases. These two tasks are almost contradictory in terms of the software requirements that they impose. On the one hand, there is a managerial and regulatory need to produce statistics that provide an overview of how certain types of failures are reduced in response to their actions. On the other hand, there is a more general requirement to identify common features amongst incident reports that should be addressed by those actions in the first place. The extraction of statistical information typically relies upon highly-typed data so that each incident can be classified as unambiguously belonging to particular categories. In contrast, the more analytical uses of incident reporting systems involve people being able to explore alternative hypotheses about the underlying causes of many failures. This, in turn, depends upon less directed forms of search. It is difficult to envisage how investigatory bodies could construct a classification scheme to reflect all of the possible causal and mitigating factors that might arise during the lifetime of a reporting system. Unfortunately, most schemes focus on the development of incident taxonomies to support statistical analysis. Relatively, few support the more open analytical activities, described above. This is reflected in the way in which most reporting systems currently rely upon relational databases.

As the name suggests, relational database techniques build on the concept of a relation. This can be thought of as a table that holds rows of similar values. For example, one table might be used to record values associated with the contributor of an incident report. The cells in the table might be used to store their name, their contact information, the date when they submitted a report and so on. Each row of the table would represent a different contributor. Of course, each contributor might make several incident reports and so another table would be used to hold this relation. The columns in this table might include the name or identifier of the person making the report, the date of the report, the plausible worst case estimate of the severity of the incident and so on. Each row in this incident table would store infromation about a different report.

Relational database offer a number of important benefits for the engineering of incident databases. One of the most important of these is the relative simplicity of the underlying concept of a relation.

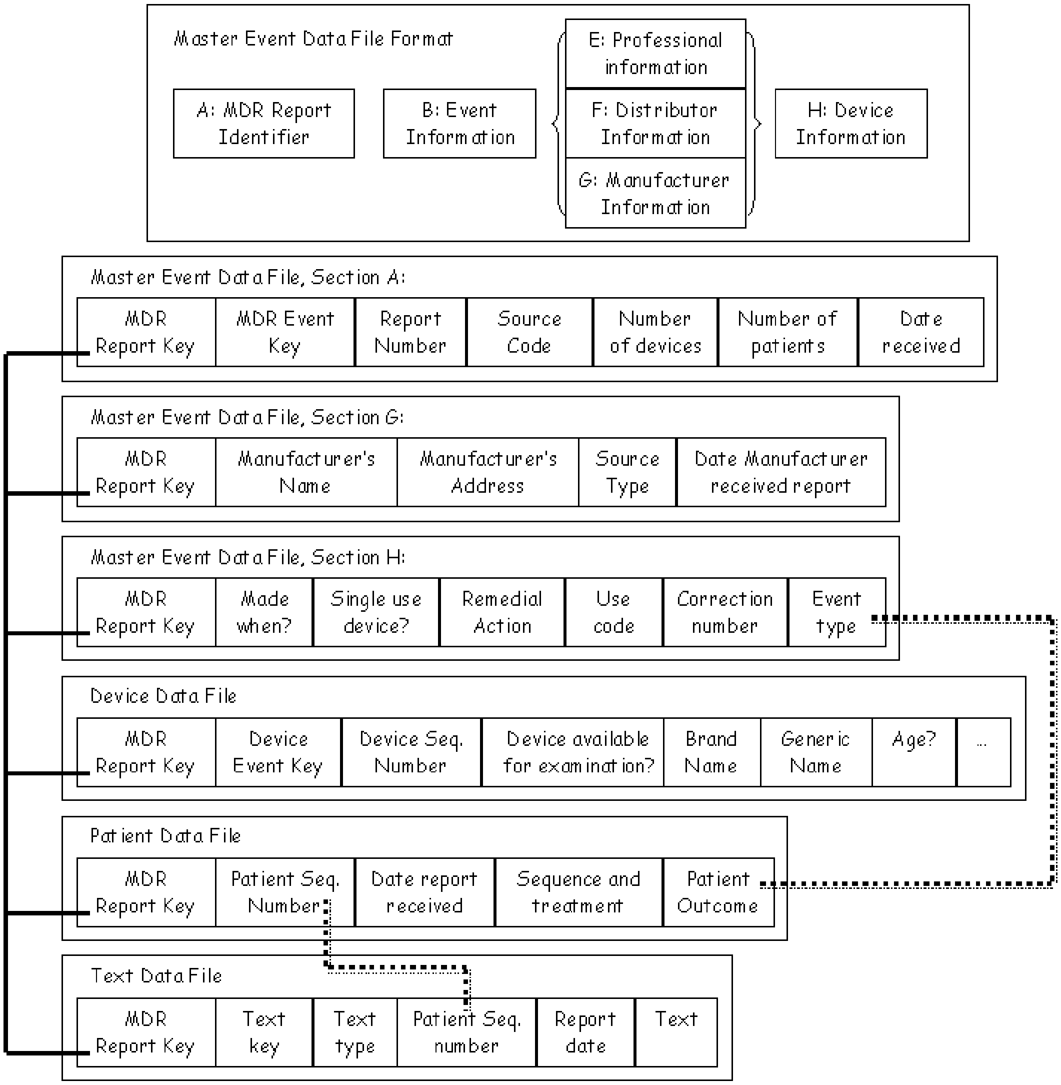

Figure 14.3: Overview of the MAUDE Relations

Unfortunately, the operational demands of many reporting systems have created the need for complex relational schemas. A schema can be thought of as a high-level model of the relationships between the different information fields that are held in the system . They are often structured in terms of objects that would be recognisable to the users of the system. Hence, as we shall see, the MAUDE schema groups information about devices, patients, manufacturers. Figure 14.3 provides a slightly simplified overview of the schema that is used to structure the FDA's MAUDE data. This overview has been reverse-engineered from the technical information that was released to enable interested parties to make use of data that is released under the US Freedom of Information provisions [268]. It is also based on a study of the way in which information is retrieved by the FDA's on-line databases. As can be seen from Figure 14.3, the data is structured around four different relations: a master event relation; a device relation; a patient relation and a text relation. For convenience, these are stored and can be retrieved over the Internet in four different files.

The Master Event Data holds information about the person or group that reports an event. This is the most complex of the relations. It distinguishes between a total of seventy-two different items of information. It also illustrates previous comments about the way in which relatively simple ideas can quickly be compromised by the operation demands of an incident reporting system. The high level structure of this file is illustrated by the top component of Figure 14.3. It is based around the idea of a nested relational schema because any entry in the Master Event Data relation is, itself, composed of more complex relations. These relations hold summary data about the nature of the event, about who has reported the event and about the devices that were involved. Section A of the meta-relation holds identification information relevant to the report. This is compulsory for all reports because, as we shall, see this acts as an index that can be used to cross-reference between the information that is held in other relations within the MAUDE system. Section B hold information about the particular event that is described in the report and again should be included for all reports. Section E is only used if the report was submitted by a healthcare professional. Similarly, Section F only applies to reports filed by device distributors. Section G only applies if the report was completed by a device manufacturer. Finally, Section H is based around a relation that is used to structure information about the devices that were involved in an incident. This should be included for all incidents. As mentioned, the precise information that is held about any particular report depends on the nature of the person or group who submitted the form. For instance, if a form was submitted by a device distributor then the master record will hold information both about the distributor and about the manufacturer that provided them with the device. In this instance, the Master Event relation would contain sections A, B, F and H. If the report is filed by a healthcare professional then they might not be in a position to enter this information into the reporting form and hence it will be omitted from the database. The relation would, therefore, be composed from sections A, B, E and H. Figure 14.3 is simplified by only showing the relations that would record a report from a manufacturer. The FDA provide summary information about the other formats [268].

It is reasonable to ask why anyone should devote this level of attention to the manner in which data is stored within an incident database. It might be argued that such details relate solely to the implementation of particular systems and are of little interest to a more general audience. Such arguments neglect the consequences that such techniques can have upon the end-users of incident databases. For instance, many local incident reporting systems adopt a more 'naive' approach and simply 'flatten out' the nested relations that we have described in the previous paragraph. This would result in every entry in the system being provided with a cell for a Health Care Professional's contact address even though the report was submitted by a distributor or manufacturer. The significance of this should be apparent if we recall that MAUDE receives between 80,000 and 85,000 submissions per year. Relatively minor changes to the relational schema can have a huge impact upon both the time that is taken to search through or download an incident database.

As mentioned, the Master-Event component of the MAUDE database is composed from nested relations. One element of this more complex structure structures the information that is necessary to unambiguously identify each incident. Figure 14.3 illustrates this by the thick line that links the MDR report key across all of the other relations in the MAUDE system. The importance of this 'key' information can be illustrated by the following example. Supposing that we wanted to find out how many patients had been injured by all of the devices produced by a particular manufacturer, we

MDR Report Key	MDR Event Key	Report Number	Source Code	No. of devices	No. of Patients	Date received
Generated	Generated	Generated	Voluntary/ User facility/ Distributor/ Manufacturer	0..Max	0..Max	Date
2339271	319405	2919016-2001-00002	Manufacturer	1	1	06/22/2001
2339103	319248	2124823-2001-00010	Manufacturer	1	1	05/20/2001

Table 14.10: The MAUDE Master-Event Relation (Section A)

could begin by using the Master Event Section G to list all of the MDR Report Keys associated with that Manufacturer's name. After having retrieved the list of MDR Report Keys we could then use Section A of the Master Event File to find out the number of patients that had been reported to be affected in initial reports to the system. It is important to note, however, that the MDR report Key is insufficient to unambiguously identify all information within the system. For instance, an incident might involve more than one device. In this case, Figure 14.3 shows how each entry in the MAUDE Device Data relation must be identified both by the MDR report key and by the Device Event Key. Again, it is important to emphasise that these implementation details have a profound impact upon the users of an incident reporting database. If they are not taken into account during the early stages of development then it can be difficult to extract critical information about previous incidents. In the example outlines above, it might be difficult to extract information about the individual devices that are involved in a single incident. If this data is not clearly distinguished then it can, in turn, become either difficult or impossible to trace the pervious performance of the manufacturers of those devices. Too often, investigators and regulators have sub-contracted the implementation of incident reporting databases with the assumption that such relational structures are both obvious and easy to implement. Equally sub-contractors have often failed to communicate the impact of these technical decisions on those who must operate incident databases. It is, therefore, hardly surprising that so many people express disappointment and frustration with the systems that they are then expected to use.

Table 14.10 provides more details about the report identification information that is held in the Master Event relation. As can be seen, each new entry is automatically assigned three reference keys: the MDR report key, the event key and a report number. The precise meaning and purpose of each of these values is difficult to infer from the FDA documentation that is provided with the MAUDE database. However, it is apparent that the MDR report key and the event key are used to index into other sources of information in the manner described above. The source code information helps to distinguish between the various groups and individuals who might submit a report to this system. Any contribution is either voluntary or is provided to meet the regulatory requirements on end-user facilities, device distributors or manufacturers. The second row of the table summarises the type of information that can be entered in each column. The third and fourth rows of table 14.10 present the values that were entered into the MAUDE database to describe two recent software related failures.

Table 14.11 characterises the nested relation inside the Master Event record that holds information about device manufacturers. Recall that this is only used if a report is submitted by a manufacturer. As can be seen, the name and address of the contributor is stored with the record. Table 14.11 is a slight simplification because the address component of this relation is itself a nested relation containing fields to store manufacturer's street name, their city and state, their telephone number and so on. It might seem like an obvious and trivial requirement to provide such a structure to record the contact information of a contributor. However, the MAUDE structure shows considerable sophistication in the manner in which this data is handled. For instance, two fields

MDR Report Key	Manufacturer's Name	Manufacturer's Address	Source Type	Date Manufact. Received
Generated	Text	Address	Other/ Foreign/ Study/ Literature/ Consumer/ Professional/ User facility/ Company rep./ Distributor/ Unknown/ Invalid data	Date
2339271	A. Maker	Somewhere	Professional, Other	05/30/2001
2339103	Another Maker	Somewhere Else	Professional, User Facility	05/18/2001

Table 14.11: The MAUDE Master-Manufacturer Relation (Section G)

are associated with street address information. Many databases simplify this into a single field and then subsequently have problems encoding information about appartments and offices that have 'unconventional' addresses. These issues are not simply important for the technical operation of the incident database. They can also have significant consequences for the running of the system. It is clearly not desirable to have investigators search for a contributor to a confidential or anonymous system in order to conduct follow-up interviews.

Table 14.11 also includes an enumerated type, or list of values, that can be used to describe the source that first alerted the device manufacturer to the incident. This information is critical and is often omitted from incident databases. Chapters 3 and Chapters 6 have argued that it is important not simply to identify the causes of a near-miss or adverse occurrence. It is also important to identify those barriers that prevented such incidents from developing into more critical failures. Any mechanism that alerts a contributor to a potential failure is an important component in the defences that protect the safety of future applications. Often this information is embedded within free-text descriptions of adverse events and it can prove to be extremely difficult to collate data about such defence. The MAUDE system avoids this problem by firstly prompting the contributor to provide this information by ticking an element in a list of the incident form and then by encoding their response within the 'source type' field of the relation illustrated in Table 14.11. The elements of this type are instructive in their diversity. Incidents may be detected from a healthcare consumer or a healthcare professional, they can also be identified by literature reviews or other forms of field study. The final rows of Table 14.11 illustrate sample values for this relation.

Table 14.12 presents the final nested component of the Master Event relation. This is completed for all submissions and provides initial details about the devices that were involved in a near miss or adverse occurrence. Again an examination of the components of this nested relation can be used to illustrate some of the points that have been made in the previous chapters of this book. For example another enumerated type is used to categorise the immediate remedial actions that a manufacturer has taken to address any incident that has been brought to their attention. This is significant because many incident and accident analysis techniques focus directly on causal events rather than examining the critical actions that are taken in the aftermath of an adverse occurrence. In this instance, the manufacturers' actions are important because they may pre-empt any further regulatory actions by the FDA, for instance, if a device recall has already been issued.

The care that has been taken in devising the FDA's relational scheme is also illustrates by the use code in Table 14.12. Figure 3.3 in Chapter 3 illustrated the higher probability of failure that

MDR Report Key	Made When?	Single use?	Remedical Action	Use Code	Correction No.	Event Type
Generated	Date	Yes/ No	Recall/ Repair/ Replace/ Relabelling/ Other/ Notification/ Inspection/ Monitor Patient / Modification/ Adjustment/ Invalid data	Initial use / Reuse/ Unknown	Previous FDA Notification Reference	Death/ Injury/ Malfunction/ Other
2339271	-	No	-	Initial use	No	Other
2339103	05/18/2001	No	Notification	Initial use	No	Other

Table 14.12: The MAUDE Master-Device Relation (Section H)

is associated during the initial period of operation for hardware systems. This occurs because of component variations but also from the problems associated with setting up devices and of learning to operate them under particular working conditions. The use code in the Master Event Relation, therefore, distinguishes between initial use and reuse of any particular device. This field also illustrates a generic problem with incident reporting databases. Ideally, we would like to provide a meaningful value for every field in every relation. This would enable use to satisfy requests of the following form 'how many incident reports related to the initial use of a device?' or 'how many reports were immediately resolved by the manufacturer issuing a recall?'. Unfortunately, lack of data can prevent investigators from entering all of the requested data into an incident database. For example, if a healthcare professional informs a manufacturer of an incident they may neglect to pass on the information that is necessary to complete the use code. In such a circumstance, the manufacturer would tick the 'unknown' category and return the form to the FDA to be entered into the MAUDE database. This would create problems because if we attempted to answer the question "how many incident reports related to the initial use of a device?' then it would be unclear how to treat these 'unknown' values. If they were excluded then this might result in a significant underestimate of the initial device set-up problems. If they were excluded then the converse problem would occur. Similar concerns can be raised about the 'Remedial Action' field in Table 14.12. In this case, the FDA analysts can enter an 'invalid data' category rather than 'unknown'. This is worrying because manufacturers might, in fact, be exploiting a range of potentially valid remedial actions that are lost to the database simply because they do not fit easily within the categories of remedial action that are encoded within this component of the Master Event relation. These concerns lead to two key heuristics for the application of relational databases to incident reporting:

- *unknown data.* If an 'unknown' value is entered for a field then a caveat must be associated with any statistics that are derived from the data in that field. Ideally, this warning should provide information about the proportion of unknown values compared to those that are known for that field.

- *invalid data.* If 'invalid data' is entered for a field then analysts should also record a reason why this option was selected. System managers should then conduct periodic reviews to ensure that important information is not being omitted through poor form design or an incomplete relational schema.

It is important to emphasise that to entirely exclude either of these categories would place severe constraints on data entry for incident reporting systems. In contrast, these heuristics are intended to ensure that relational schemas continue to offer the flexibility that is necessary when encoding incomplete accounts of adverse occurrences and near misses. They are also intended to ensure that this flexibility does not compromise the integrity of the information that is derived from incident databases [224].

MDR Report Key	Device Event Key	Device Seq. No.	Device Available?	Age	Brand Name	Generic Name	Baseline id	...
Generated	Generated	1.. Max	Yes/ No/ Returned/ No answer	0..Max	Text	Text	Generated	...
2339271	328578	1	No	2	The Item	HNID Panel	K833027	...
2339103	328407	1	No	3	Product	Central Station	K954629	...

Table 14.13: The MAUDE Device Records

Table 14.13 illustrates how MAUDE holds further device information in a separate relation. This separation can be explained by the observation that the Master Event relation holds information that is derived from the incident report. The device relation, in contrast, can hold information that need not be available from the initial report. For instance, Table 14.13 includes a field that is intended to hold information about any baseline report that is associated with a device. Chapter 6 has described how baseline reports must be submitted in response to the first reportable incident involving a particular device. It provides basic device identification information including: brand name, device family designation, model number, catalogue number and any other device identification number. This information helps ensure clear, unambiguous device identification. From this it follows that if the incident described in the Master Event relation is not the first occurrence to affect a device than the baseline report summarised in the device relation will be based on previous information. It is again important to emphasise that the relation shown in Table 14.13 simplifies the data that is actually held by the MAUDE system. The device relation ic composed of 43 individual fields. These include information not only about the particular device that was involved in the incident but also about the product range or family that the device belongs to. Such details again emphasise the importance of considering each element of a relational scheme in order to ensure that it captures all of the information that may subsequently help to identify patterns of failure in similar devices.

Table 14.13 illustrates a number of further, generic issues that affect relational schemas in many different incident databases. For instance, the device sequence number helps to distinguish between the different items of equipment that can be involves in any single incident. In order to refer to any particular device record, therefore, it may be necessary to supply both the MDR report key and the sequence number. Alternatively, a device event key can also be supplied to unambiguously identify a device record. Table 14.13 also captures some forensic information, including whether or not a particular device is available for examination. As with the use code in Table 14.12, this field also permits the entry of an unknown value. In this case it is termed ‘no answer’. This illustrates a potential problem for the coders who enter the data into the system. They must be trained to distinguish between, or conversely to ignore, the subtle differences in terminology that are used to represent null values in these two different contexts.

A further relation holds information about the patents that were affected by a near miss or adverse occurrence. This is completed even if there were no long term consequences for the individuals who were involved in an incident. Just as there can be several devices that are involved in an incident, there can also be more than one patient. In consequence, any individual patient record must be identified both by the MDR report key and also by the patient sequence number. The sequence numbers that are associated with any incident can be inferred from Section A of the Master Event relation because this records the total number of patients that were affected by an incident. The integrity of the database therefore depends upon a number of assumptions:

MDR Report Key	Patient Sequence No.	Date Received	Sequence Treatment	Patient Outcome
Generated	0..Max	Date	Sequence-Treatment pair	Life threatening/ Hospitalization/ Disability/ Congenital Abnormality/ Requireed Intervention/ Other/ Unknown/ No information / Not applicable / Death / Invalid data
2339271	1	06/22/2001	-	Other
2339103	1	05/20/2001	-	Other

Table 14.14: The MAUDE Patient Records

1. the Master Event record must accurately record the total number of patients that were affected in an incident.
2. a different Patient Data record must be stored for each patient involved in an incident.
3. each Patient Data record must include a unique Patient Sequence Number and these must follow consecutively from 1 to the total number of patients stored in the Master Event record with the same MDR report key.

If any of these integrity constraints are violated then there is no guarantee that it will be possible for an implementation of the database to return the patient records of all individuals who may have been affected by a near miss or adverse occurrence. For instance, if a patient record was allocated a sequence number greater than the maximum number of patients noted in the Master Event record then doubts would be raised about the reliability of that data. It is also likely that the algorithms for assembling information about an incident might miss the additional patient record if they assumed that the Master Event record was correct.

There are a number of similar constraints that must be observed by those who maintain the MAUDE system. For example, the device sequence number might be related to the total number of devices in the same manner that the patient sequence number is related to the total number of patients. There are further examples. For instance, the Master Event relation, Section H, contains information about the nature of the adverse event. Analysts must enter whether the incident resulted in a death, injury, a malfunction or some other outcome. Similarly, the individual patient records include a 'patient outcome' field that distinguishes between the following categories: life threatening; hospitalization; disability; congenital abnormality; requireed intervention; other; unknown; no information; not applicable, death and invalid data. Clearly the integrity of the database would be compromised if a patient record indicated that a fatality was associated with a particular MDR report key while the Master Event relation showed that the outcome was a malfunction. Fortunately, many database management systems provide explicit support for automating these consistency constraints. They will alert users to potential problems if they arise during data entry. It is, however, less easy for these systems to help users distinguish between the overlapping categories that the FDA have introduced for some fields. These include the 'other', 'unknown', 'no information', 'not applicable' and 'invalid data' options, mentioned above. Later sections will describe the problems that coders have experienced in choosing between these different values when they complete report forms and enter them into incident databases.

Table 14.14 also includes information about the treatment that the patient received following an incident. Any individual patient may receive a number of different treatments. In consequence, the

MAUDE relation includes a sequence-treatment pair. This simply associates a number with each of the treatments that was used on that particular individual. It would be possible to construct a further relation that holds more detailed information about each treatment. This could be indexed by the MDR report key, the patient sequence number and the sequence number of the treatment. The data that is released by the FDA from the MAUDE system does not do this. Instead it adopts this compromise approach that resembles a compound attribute [224]. The more general point here, however, is that the database records remedial actions that relate to individual devices, such as product recalls, and the treatments that are taken to counter any adverse consequences for an incident to the patients that are affected. In other words, MAUDE illustrates the broad approach that must be taken when considering what information to capture about the response to any incident.

Table 14.15 provides an overview of the final relation that is used to structure incident data in the FDA's MAUDE system. This relation os central to the success of the system and it represents a solution to a generic problem that affects all incident databases. The previous relations have provided a means of grouping or structuring related information. The patient relation holds information about an individual patient, the device record holds information about an item of equipment that is implicated in an incident and so on. Each element of information that might be placed within a field in one of these relations has an associated type. Most of these types are constrained. For instance, the event type in Section H of the Master Event record can only take the values: death; injury; malfunction or other. This helps to reduce coding problems. Analysts must only differentiate between a few values rather than the subtle differences that might exist between a larger range of potential values. They also help to provide numerical results for statistical analysis. It is relatively easy to sum the total number of incidents which were classified as resulting in a death. This would be far harder if analysts were able to enter any free text value that they liked in the event type field. Analysis might use the terms 'fatal' or 'fatality', 'dead' or 'death' and so on. An automated system would then have to predict all of these potential values and recognise that they were equivalent in calculating any summary statistics. These problems are avoided by have a small range of admissible values that are associated with the various fields in a relation schema. Similarly, the fields in the schema also define a minimum data-set that should be obtained about each incident. The previous paragraphs have described how this minimum data-set can depend upon the nature of the incident report. The information that is available for voluntary reports by a healthcare professional might be very different than that which is available following a mandatory report from a manufacturer. Similarly, we have also described how lack of evidence in the aftermath of an incident can prevent investigators from satisfying the minimum requirement implied by a relational schema. The key point is, however, that by providing 'invalid data' or 'unavailable' options in the database, investigators can be sure that this analysts were prompted for this information when they entered incident data into the system. Any omission, in principle, should be due to the constraints that characterise the aftermath of the incident rather than neglect on the part of the analyst.

Unfortunately, the strengths that the relational model derives from the explicit grouping of related fields of typed information can also be a significant weakness for many incident reporting systems. As we have seen, many near misses and adverse occurrences cannot easily be characterised into the relatively small number of fields that have been introduced in the previous pages. For example, the MAUDE relational schema offers almost no opportunity for analysts to enter the contextual information about workplace factors, such as time pressure or staffing issues, that have been stressed in previous chapters. If any database only recorded the typed information mentioned above then subsequent investigators would derive a very biased view of the causes of previous incidents. In consequence. most relational systems also provide for storage and retrieval of large textual accounts. For instance, Table 14.15 shows how there are two different types of text that can be associated with each MDR report key. Event description summarise any additional information about the immediate course of a near miss or adverse occurrence that cannot be provided in the previous fields. The manufacturer narrative, in contrast, provides an opportunity for the producers of a device to respond to any incident reports. As can be seen in Table 14.15, this response can include information about subsequent studies into the cause of an incident. Such studies must describe the methods used and the results that were obtained. It might, therefore, be argued that a nested relation could be used to distinguish these approaches. This would enable analysts to

MDR Report Key	Text Key	Text Type	Patient Seq. No.	Report Date	Text
Generated	Generated	Event description/ Manufacturer narrative	0..Max	Date	Text
2339271	1173556	Event description	1	06/22/2001	User reported a clinical isolate was identified on Microscan HNID panel read by the walk-away instrument system as Neisseria Gonorrhoeae with 99% probability. The specimen was from a blood source from PT. due to the unusual source for this organism the specimen was sent to the State Health Dept Reference Lab for confirmation. The State Reference Lab Identified the organism as Neisseria Meningitis.
2339271	1173558	Manufacturer narrative	1	06/22/2001	H.6 EVAL METHOD: obtained clinical isolate from customer and tested on products involved i.e. HNID panels and Microscan Walk-away Instrument system. Reviewed complaint history, performance evals, labeling and literature regarding reported issue. H.6. RESULTS: Biotype reported by user was duplicated by Microscan Technical Services lab. Atypical results suggest and footnotes indicate N. Gonorrhoeae identification required additional tests to confirm. Results from add'l tests should lead to a presumptive identification on N. Meningitis. Results of complaint history review revealed a very low complaint volume for this issue...

Table 14.15: The MAUDE Text Records

pose queries about which manufacturers had used a particular evaluation method in response to an incident failure. This is not, however, possible using MAUDE because a general text field is used rather than the more strongly typed approaches that are embodied in previous relations.

The examples in Table 14.15 illustrate the way in which several textual accounts can be associated with a single incident. As mentioned, there is both an event description and a manufacturer narrative for event report 339271. It is also important to realise that more than one individual may be affected by an incident. In such circumstances, an event description can be associated with each person who was, or might have been, injured. The previous table, therefore, includes the patient sequence number associated with each text report. For this it follows that in order to uniquely identify any particular report, analysts will have to supply the MDR report key, the text type and the patient sequence number. In some cases, manufacturers may make more than one response to an incident. If such multiple responses were admitted then analysts would also have to specify the date of the message that they were interested in retrieving. Such requirements appear to introduce unnecessary complexity into an incident database. It is important to remember, however, that there could be profound implications if it appeared that a subsequent response to an incident had in fact been made in the immediate aftermath of an adverse report. Unless the database supports such version control, investigators would have no means of knowing when a narrative was introduced into the system.

This section has provided a relatively detailed analysis of the relational model that is used by the FDA to structure the data contained in their Manufacturer and User Facility Device Experience Database. The level of detail in this analysis is justified by the observation that many reporting databases have failed to provide their expected benefits precisely because those who have commissioned these systems have failed to pay sufficient interest to these details. Conversely, the sub-contractors who are typically enlisted to implement these systems often fail to explain the consequences of particular relational schema both on the queries that can be posed of the system and on the performance that can be obtained as the size of the system grows. Our analysis has also helped to identify a range of benefits that can be derived through an appropriate use of the relational model that is embodied in most incident databases:

- *Analytical help in developing the relational schema.* It can be argued that the process of developing the relational schemas that underly many databases can help to indentify key information requirements. This process is supported by a range of well-documented methods, including entity-relationship modelling and the analysis of normal forms [224]. In particular, these approaches can help to expose the integrity constraints that must be satisfied by any implementation. Although these techniques have their limitations and none are specifically intended to support the development of incident databases, they do have the strong advantage that they are 'industry standard' and hence widely understood. This introduces an important paradox because the reverse enginering of many incident databases has revealed important structural weaknesses, which suggest that many of these systems have been built without the benefit of these relatively simple engineering techniques [414].

- *Analytical help in guiding incident classification.* The development of a relational schema is intended to enable investigators, regulators and safety managers to classify incident reports so that they can be analysed and retrieved at a later date. The process of constructing a relational schema, therefore, forces people to consider the forms of analysis that any system must support. This leads to an important decision. Either the person submitting a form must indicate appropriate values from an incident classification or the analysts must codify a less structured account into the fields that are included in a relational schema or a hybrid model can be adopted where the contributor performs a 'first pass' classification that is then refined by the investigator. No matter which approach is adopted, the key point is that the development of the incident database must have an impact on the manner in which data is both elicited and codified. It is, therefore, extremely difficult to simply bolt-on an existing database to an incident reporting system where either contributors or analysts will have to adjust their behaviour to support the values that are built into a relational schema. In such circumstances, rather than guiding analysts and contributors towards an appropriate classification they can find that a database forces them to 'squeeze' or 'massage' an incident into inappropriate data structures.

- *Efficiency.* One of the key technical benefits behind the relational approach is that it helps to avoid the duplication of redundant information. Ideally, we might store information about a device manufacturer once. Similarly, an optimised system would only ever store a single record about any particular device. This would record the complete service and version history of that item. A link could then be made from an incident report to a device record and from there to the associated manufacturer. An alternative model would be to duplicate manufacturer information each time a new incident record was created. This is not only wasteful in terms of the storage that is required, it can also significantly increase the amount of time that is required to collate incident information. The technical reasons for this relate to the search latencies that are associated with primary and secondary storage. Relational techniques can use indexing so that once a common item of information is stored in main memory then those details do not then need to be repeatedly fetched from slower secondary media [224].

This is a partial summary, however, it is also important to stress that our analysis has identified a number of problems with the use of relational databases for incident reporting. For instance, many of these applications rely upon strong typing to clearly distinguish between the admissible values that can be entered into each field. This creates problems because in the early stages of an investigation it is often impossible to be certain about which values might hold. A good example of this might be the problems associated with any assessment of the consequences of an incident based on a clinical prognosis. This uncertainty results in a proliferation of 'unknown' values that make it very difficult to interpret the accuracy of statistics that are derived from incident databases. Similarly, it can be very difficult for analysts to accurately and consistently distinguish between the numerous values that might be entered into particular fields within a relation. This can result in similar incidents being classified in a number of different ways within the same relational schema. It can also lead to 'not applicable' values being used as a default. The previous discussion has also identified the potential vulnerabilities that can arise from the relationships that often exist between the components of a schema. In particular, problems can arise from the way in which MAUDE links the maximum number of patients and devices in the Master Event record to provide a range for patient and device sequence numbers. Automated support must be provided to ensure that consistency requirements between these linked values are maintained throughout the lifetime of the database. This is a partial summary. the following sections expand on the problems that can affect the use of the relational model for incident reporting databases. Subsequent sections then go on to review further computational techniques that can be used either to replace or augment this approach.

Problems of Query Formation

Previous sections have described how the relational model can be used to reduce the storage requirements and increase the speed of queries that are performed on incident databases. It provides further advantages. For instance, search requests can be formulated a using relational algebra. The operators within these languages have a close relationship to the operators of set theory, such as union, intersection and set difference. This offers a number of benefits. Firstly, the components of the relational algebra should have a clear semantics or meaning. Users can apply the basic ideas in set theory to gain some understanding of the query languages that are supported by most relational databases. There are further benefits. As most implementations exploit set theoretic ideas, it is therefore possible to apply knowledge gained from one relational database system to help understand another. The mathematical underpinning of the approach support skill transfer and a certain degree of vendor independence. Unfortunately, as we shall see, relatively few investigators or safety managers have acquired the requisite understanding of set theory or of relational algebra to exploit these potential benefits. The following sections provide a brief overview of the relational operators applied to an incident database. The intention is both to illustrate the potential application of this approach and also to illustrate some of the complexity that can arise from the relational algebra.

Before discussing the set operators, mentioned above, it is necessary to introduce two additional elements of the relational algebra: SELECT and PROJECT. The SELECT operator is usually represented in the relational algebra by σ and is applied in the following manner:

$$\sigma < \mathit{selection\_condition} > (\mathit{Relation}) \tag{14.1}$$

The selection condition is a Boolean expression that is usually formed from attribute names, operators and constants. The operators include $=, >, \leq, <, \geq$. Attribute names denote particular fields in a relation. For instance, Table 14.12 includes the attributes MDR report key, made when? single use? remedial action, use code, correction number and event type. The constant values include elements of the classification scheme that might be entered into these fields. For example, the Use Code constants include 'Initial use', 'Reuse' and 'Unknown'. We can put all of this together in the following manner:

$$\sigma < \mathit{Use\_Code} = \mathit{Initial\_Use} > (\mathit{Table}\ 14.12) \tag{14.2}$$

This expression would yield all of the entries in Table 14.12 which were associated with the initial use of a device. One of the benefits of this approach is that we can combine elements of the algebra to form more complex expressions. For instance, we might want to SELECT all entries that relate to either the initial use or reuse of medical devices:

$$\sigma < \mathit{Use\_Code} = \mathit{Initial\_Use} > \mathit{OR} < \mathit{Use\_Code} = \mathit{Reuse} > (\mathit{Table}\ 14.12) \tag{14.3}$$

The SELECT operation can be thought of as selecting a row from one of the relations in a database schema [224]. The PROJECT operator, in contrast, can be thought of as selecting particular columns within a relation. The application of this operation can be illustrated as follows. The first sentence denotes the general form of the PROJECT operator. The second sentence shows how it can be applied to the relation in Table 14.12. This would yield a list of all MDR keys together with the date when the corresponding device was manufactured and the type of event it was involved in:

$$\pi < \mathit{attribute\_list} > (\mathit{Relation}) \tag{14.4}$$

$$\pi < \mathit{MDR\_report\_key}, \mathit{made\_when?}, \mathit{event\_type} > (\mathit{Table}\ 14.12) \tag{14.5}$$

There are a number of important features of the SELECT and PROJECT operators. For example, if the projection is applied to non-key fields then it is likely that it will yield duplicate values. For example, if we omitted the MDR Key field in the previous example, we might derive a number of reports in which the devices were made on the same day and produced the same outcome. The PROJECT operation filters these duplicate values because the result must itself be an operation. Brevity prevents any sustained analysis of these more detailed features, the interested in reader is directed to [224]. In contrast, the following paragraphs focus on the main features of the relational algebra. The intention is to illustrate both the power of the approach but also the usability problems that can prevent many investigators from exploiting this language as a means of interacting within incident databases. It is often necessary to combine the operations in the relational algebra to form more complex requests. For example, we might wish to create a list of the dates and MDR keys for all incidents that occurred during the initial use of a device. This can be done in the following manner. Both of the following forms are equivalent, however, the operations (14.7) and (14.8) make use of the RENAME (←) operator to hold the intermediate result of the SELECT operation. This can provide important benefits as users form more complex queries:

$$\pi < \mathit{MDR\_report\_key}, \mathit{made\_when?} > (\sigma < \mathit{Use\_Code} = \mathit{Initial\_Use} > (\mathit{Table}\ 14.12)) \tag{14.6}$$

$$\mathit{TEMP\_RELATION} \leftarrow (\sigma < \mathit{Use\_Code} = \mathit{Initial\_Use} > (\mathit{Table}\ 14.12) \tag{14.7}$$

$$\pi < \mathit{MDR\_report\_key}, \mathit{made\_when?} > (\mathit{TEMP\_RELATION}) \tag{14.8}$$

As mentioned, an important strength of the relational algebra is that it builds upon the relatively well-known concepts of set theory. For example, the UNION operator can be used to produce a relation that is composed from the set of tuples that are in one or other or both component relations. This can be illustrated by the following example. Suppose that an investigators wanted to derive the MDR event keys for all incidents involving devices that were manufactured on or before 1996 or that were in use for more than six years. This can be done in three stages. Firstly, (14.9) and (14.10) identify the devices that were manufactured during or before 1995. Then (14.11) and (14.12) extract those devices that have been in operation for six or more years. Finally, (14.13) forms the union of the two previous stages of the operation:

$$Temp1 \leftarrow \sigma < made\_when? \leq 1995 > (Table\ 14.12) \tag{14.9}$$

$$Old\_devices \leftarrow \pi < MDR\_Report\_Key > Temp1 \tag{14.10}$$

$$Temp2 \leftarrow \sigma < age? \geq 6 > (Table\ 14.13) \tag{14.11}$$

$$Old\_models \leftarrow \pi < MDR\_Report\_Key > Temp2 \tag{14.12}$$

$$All\_Old\_Devices \leftarrow Old\_models \bigcup Old\_devices \tag{14.13}$$

$$\tag{14.14}$$

INTERSECTION and SET DIFFERENCE can be used in a similar fashion. For example, we could identify the MDR Report Keys of devices that were manufactured during or before 1996 but which have not been in operation for 6 or more years using SET DIFFERENCE in the following manner:

$$Legacy\_Devices \leftarrow Old\_models - Old\_devices \tag{14.15}$$

Conversely, we can identify those models that were manufactured during or before 1996 and which have been in operation for six years or more using INTERSECTION. This illustrates the flexibility of the relational algebra. It is possible to use the set theoretic operators to express a range of relatively complex constraints that can be applied to relatively simple relational schemas:

$$Obsolete\_Devices \leftarrow Old\_models \bigcap Old\_devices \tag{14.16}$$

The JOIN ($\bowtie$) operation is used to combine related tuples from two relations to form a single relation. If the first relation has N attributes and the second relation has M attributes then the resulting relation will have N+M attributes. The application of this operation can be illustrated in the following manner. The first sentence illustrates the general form of the JOIN relation. The second sentence illustrates how it can be used together with the *All_Old_Devices* that was derived from (14.13) to extract all of the patient related information for incidents that involved devices, which were either manufactured during or before 1996 or that have been in use for more than six years:

$$Relation\_1 \bowtie < join\_condition > Relation\_2 \tag{14.17}$$

$$Patients\_affected \leftarrow All\_Old\_Devices \bowtie < MDR\_Report\_Key > (Table\ 14.14) \tag{14.18}$$

As before, a number of additional details must be considered when using this operator. For instance, the results of a JOIN operation do not typically include any tuples with NULL values in the parameters of a *join_condition*. This property of the relational algebra can have important consequences for incident reporting databases. The proliferation of null values can mask a large number of candidate incidents that might have been included within the results of a particular query if more information had been obtained about the adverse occurrences that they document. Hence, the results of queries that use the JOIN operator may significantly under-report all possible candidate incident records. They may also mask the number of reports that are omitted by simply 'dropping' all candidate tuples with NULL values. These problems cannot arise in the previous example because we have joined the relations on the primary key that is used throughout the MAUDE system, in other words the MDR Report Key cannot contain a null value. Problems would, however, arise if we attempted to perform a JOIN using the Patient Outcome or Source Type fields. As we shall see, MAUDE avoids this problem by constraining the queries that can be performed using these relations. In general, this problem can be avoided by using a form of the OUTER JOIN operation. Unfortunately, such distinctions are often not apparent to those who must learn to use incident databases [472].

A number of pragmatic observations can be made about the relational algebra and its use within incident reporting systems. For those with a mathematical background, it offers a clear semantics and its origins in set theory can reduce training times. For those without such a training, it can appear to be both complex and confusing. There are also other aspects of the language that often irritate both sets of users. A particular feature is that the user must specify a precise ordering for each operation. If they are performed in any other sequence then the result may not be what the user had anticipated. The syntax associated with the relational algebra has also been criticised as opaque and difficult to learn [224]. It is for this reason that the Structured Query Language (SQL) has emerged as a standard means of interacting with many database systems. There are strong differences between the elements of this language and the relational algebra. Duplicate values are allowed within in SQL tables, they are not permitted within relations. Hence the mathematical underpinnings of SQL are based more on bags that sets of tuples. There are further differences. For instance, the SQL SELECT statement has no formal relationship to its counterpart that was introduced in previous paragraphs:

```
SELECT <attribute_list>
FROM <table_list>
WHERE <condition>
```

In this general form, `<attribute_list>` is a list of the attribute names or fields whose values are to be retrieved by the query. A `<table_list>` is a list of the relation names that will be processed by a query. As might be expected, `<condition>` is a boolean expression that identifies those relation components that are to be retrieved. The application of this form can be illustrated by the following query, which retrieved the MDR Report Keys and the date of manufacture for devices that were either being used for the first time or were being re-used when an incident occurred:

```
SELECT <MDR_Report_Key, made_when?>
FROM Table.14.12
WHERE <Use_Code= Initial_Use> OR <Use_Code= Reuse>
```

The SQL query is intended to be declarative. Users should not have to worry about the precise ordering of the individual terms in an expression. It can be contrasted with the corresponding formulae in the relational algebra in which the select must be performed before the projection because the projection would strip out the information about the use code that is used as the basis for the selection operation:

$$\pi < \mathit{MDR\_report\_key}, \mathit{made\_when?} > (\sigma < \mathit{Use\_Code} = \mathit{Initial\_Use} > (\mathit{Table}\ 14.12)) \quad (14.19)$$

SQL offers further benefits. In particular, it is possible to construct complex queries that select elements from several different relations or tables. The following example extracts the MDR report key, any remedial actions and the patient outcome for any incidents involving the CIU panel brand of devices. Notice that the patient outcome is derived from the data that is held in the MAUDE patient file, illustrated by Table 14.14, while the remedial action is associated with a particular device report, illustrated by Table 14.12:

```
SELECT <Table.14.12.MDR_Report_Key, Table.14.12.Remedial_action,
          Table.14.14.Patient_outcome>
FROM Table 14.12,Table.14.14
WHERE <(Table.14.12.MDR_Report_Key = Table.14.14.MDR_Report_Key) AND
          (Table.14.12.Brand_name = ‘CIU Panel’)>
```

In addition to these mechanisms for complex query formation, SQL also offers a limited range of statistical operations that can be used to support the monitoring functions, which will be discussed in Chapter 15. These functions include SUM, MAX, MIN and AVG. These can be applied to a set or bag of numeric attributes. The following example illustrates a query to find out the total number of patients that have been affected by adverse incidents, the maximum number affected by a single incident and the average number of patients affected:

```
SELECT <SUM(Number_of_patients), MAX(Number_of_patients), AVG(Number_of_patients)>
FROM Table 14.10
```

The rapid development of SQL as a standard for interaction with relational databases provides a strong indication of its advantages over the 'raw' relational algebra that has been illustrated in previous pages. However, these benefits do not provide a panacea for the implementation of incident databases. Many of the advantages that SQL offers can only be appreciated by programmers and developers who have the necessary background to exploit many of its more advanced features. The majority of safety managers, of regulators and of incident investigators have little appreciation of how to use SQL. This creates considerable practical problems. For instance, the '*query paradox*' arises because those people who can best exploit data about previous failures lack the technical expertise to form the queries that reveal hidden patterns within the data. In contrast, the individuals who have the technical expertise to form appropriate queries often lack the understanding of the application domain that is necessary to identify what questions to ask the incident databases. This paradox does not simply affect incident databases. Its consequences are, however, potentially more serious given the nature of the data that is held in this systems. There are a number of potential solutions. For instance, database experts and information technologists might be trained to have a greater appreciation of the application domain. Alternatively, investigators, regulators and safety managers might be trained to have a greater appreciation of the technical underpinnings of the systems that they use to support the everyday tasks. Unfortunately, neither of these alternatives seems to have been followed in any systematic manner. It is more common to find a lack of understanding between those who maintain reporting databases and those who must use them [423]. This is often revealed in complaints that the system will not provide access to data that the users believe it 'must hold'. Conversely, administrators are often faced with demands to support facilities that cannot easily be provided using relational databases, such as free-text retrieval.

The tensions that are created by the query paradox focus on the design of the user interface to incident databases. These, typically, attempt to simplify the task of interacting with large collections of data by supporting a limited number of pre-formulated queries. All that the use has to do is specify values for the particular fields that they are interested in. This can create potential conflicts because these pre-canned queries are unlikely to satisfy the diverse requirements of many potential users. Previous paragraphs have emphasised the difficulty of predicting all of the possible queries that investigators and regulators might want to pose to such a system. Phrases such as 'data mining' and 'exploratory analysis' are often used to publicise these systems but these activities are hardly supported by the limited numbers of 'pre-canned' queries that are supported by many incident databases. In consequence, support staff are often faced with continual requests to perform one-off analyses that cannot easily be constructed from the existing interface.

The top image of Figure 14.4 illustrates the screen that provides web-based access to the MAUDE system. The bottom of the two images shows the list of incidents that can be derived from a particular query. If the user selects any one these 'hits' they can view a relatively complete summary of the various fields that have been described in previous paragraphs and which are illustrated in Tables 14.10, 14.11, 14.12, 14.13, 14.14 and 14.15. As can be seen from the top screen, users can either enter specific values for a relatively small number of fields or they can perform a free text search. Subsequent sections will describe the strengths and weaknesses of such 'free-text' searches in greater detail. For now it is sufficient to observe that the standard query interface only provides a very small subset of the potential queries that users might pose of the FDA's data. Recall that the main data file holds a relation with 72 fields. The device file relation provides a further 43 fields. It would be difficult to design a graphical user interface that would enable a untrained user to form the wide range of SQL queries that might be performed on such data sources. The relative simplicity

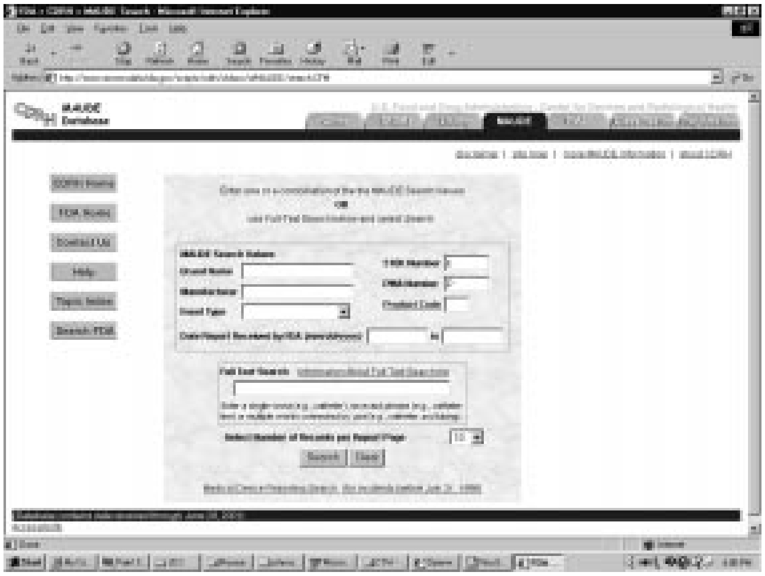

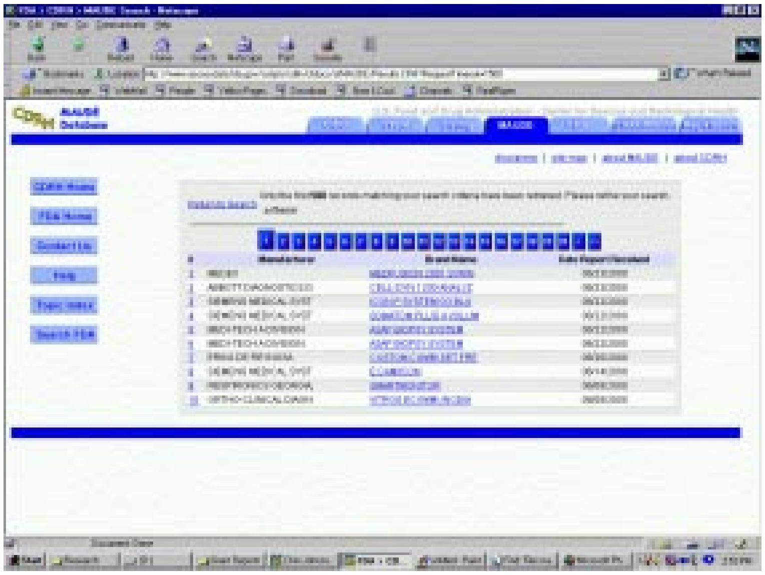

#	Manufacturer	Brand Name	Date Report Received
1	[illegible]	[illegible]	[illegible]
2	ABBOTT DIAGNOSTICS	[illegible]	[illegible]
3	SIEMENS MEDICAL SYST	[illegible]	[illegible]
4	SIEMENS MEDICAL SYST	[illegible]	[illegible]
5	MEDI-TECH A DIVISION	[illegible]	[illegible]
6	MEDI-TECH A DIVISION	[illegible]	[illegible]
7	[illegible]	[illegible]	[illegible]
8	SIEMENS MEDICAL SYST	[illegible]	[illegible]
9	[illegible]	[illegible]	[illegible]
10	ORTHO-CLINICAL DIAGN	[illegible]	[illegible]

Figure 14.4: The MAUDE User Interface

of the top screen in Figure 14.4 can, therefore, be seen to have strong design strengths. It does not daunt initial users with a vast array of bewildering options. Equally, however, it does not support the more sustained analysis of trends and patterns that might be performed by users who have a more extensive knowledge of the technical under-pinnings of relational databases.

It is important to emphasise the complex nature of the paradox that was introduced in previous paragraphs. Even if investigators are trained to appreciate the concepts and mechanisms of relational databases there are many potential further pitfalls. Many professional software engineers fail to construct 'correct' queries using relational query languages such as SQL [702]. In other words, they rely on queries that will not return the information that they are believed to. There are technical issues, such as the different semantics associated with the JOIN operator, that complicate the application of these techniques. As we have seen, this is a particular problem for incident databases that are likely to contain many NULL values in the aftermath of a near miss or adverse occurrence. These technical issues are, arguably, less significant than the problems of ensuring that incident data is correctly entered into the system in the first place.

Problems of Classification

The previous paragraphs have described how relational databases are constructed around a number of fields. Each of these fields captures particular values. For instance, the event type in Section H of the MAUDE Master Event record can be: death; injury; malfunction or 'other'. The use of such schemas offers numerous benefits. For example, it provides a national framework for the collection and analysis of incident data. Any organisation that contributes to the system must provide their data in a format that can easily be integrated into the relational schema. It is difficult to under-emphasise the practical importance of this. Other industries, have experienced considerable difficulties in exploiting incident data precisely because they lack an agreed format that can be used to structure incident information. For example, Boeing currently receives data about maintenance incidents from many customer organisations. Each of these organisations exploits a different model for the records in their relational systems. As a result, the aircraft manufacturer must attempt to unify these ad hoc models into a coherent database. The GAIN initiative has taken considerable steps to address this problem [308]. At present, however, it can be difficult to distinguish between bolts that have failed through design flaws and bolts that have failed because of over-torquing by maintenance engineers. Sam Lainoff summarised the problems of populating relational databases:

> "There is no uniform reporting language amongst the airlines, so it's not unusual to find ten different ways of referring to the same thing. This often makes the searching and sorting task a difficult proposition? The data we have won't usually permit us to create more refinement in our error typing. But at times it will give us enough clues to separate quality problems, and real human error from pure hardware faults." [472].

The previous quotation blurs the notion of a taxonomy or language for incident reporting and the relational technology that is used to retrieve incident reports. This confusion is understandable. Relational databases implement data models. The popularity of this technology has created a situation in which it is difficult to envisage the development of a taxonomy or data model without some corresponding computer-based implementation. On the other hand, there have been a number of national and international initiatives to develop taxonomies for incidents and accidents without considering the need for tool support. The US National Patient Safety Foundation (NPSF) clarifies some of these important distinctions in its analysis of the Aviation Safety Reporting System (ASRS) . The final sentence is particularly important because it identifies the opposite problem to that faced by Boeing. The previous quotation illustrates the difficulty of synthesising incident data that does not share a common model or taxonomy. The following quote recognises the danger that events will be analysed until they fit the taxonomy even if that taxonomy does not accurately represent the incident under consideration:

> "In the discussion about incident reporting, it was pointed out that the ASRS uses an extensive indexing system, but this is used to collect related subsets of narrative cases from the database that pertain to a theme or question. The indexing system does not

> work automatically but is a tool used by the staff to carry out analyses and to assist outside parties use the database in their analyses. The indexing is used as a tool in analysis; the classification system it represents is not the analysis."

The importance of an appropriate data model cannot be under-estimated. For example, these taxonomies drive the statistical analyses that are often cited as a primary benefit of incident reporting. We have already described how the event type in Section H of the MAUDE Master Event records whether an incident involved a death, injury, malfunction or 'other'. Summing the number of records in each category can provide valuable information about the types of occurrences and near-miss incidents that are reported to the system. Equally, if we were concerned to identify the number of incidents that resulted in a particular type of injury then we must look to other fields in the MAUDE system. If none of them satisfied our information requirement then we might not be able to report on those types of incident.

The incident models that are embedded within relational databases have further benefits. They help to guide the local analysis and classification of incident reports. This is a significant benefit for large-scale systems. Central organisations may lack the necessary local insight to drive the classification of a particular incident. They may also lack the resources that are required to centralise the analysis of every potential report. By devolving the classification process to regional or local representatives, central investigators can focus on responding to higher-criticality incidents or to those exceptional incidents that do not fit into the existing taxonomy. A key point here is that the values in the data model provide powerful guidance to those individuals who are documenting an adverse occurrence. They provide a prompt for the type of information that must be provided. They also indicate the particular values that each item of information might take. It is, therefore, often argued that the data models that are embeded within relational databases help to improve inter-rater reliability during the analysis of incident reports.

It can be extremely difficult to construct a taxonomy that is capable of capturing all of the information that people might want to extract about adverse occurrences and near-miss incidents. For instance, one approach is to rely upon a small number of high-level categories for most of the data fields. Information would be gathered about 'software failures' rather than 'floating point exceptions'. Similarly, incidents might be characterised by 'human error' rather than 'poor situation awareness'. This approach has the advantage that many high-level categories will be resilient to change. The particular forms of software failure that may be reported to a system can be affected by changes in the underlying technology. Similarly, the findings of human factors research are likely to have a more profound impact on detailed distinctions than they are upon broader categories. There are further benefits. For instance, by restricting the number of distinctions that must be made between different types of incident data it is possible both to increase inter-rater reliability and reduce potential training times for local and regional analysts.

Unfortunately, high-level taxonomies suffer from a number of limitations. The elements of these models often fail to capture the particular details that characterise many incidents. They may, therefore, omit information that might contribute to the safety of future systems. In particular, high-level taxonomies tend to support retrieval systems that yield very low precision for many of the queries that users want to pose. For instance, a request for information about 'floating point exceptions' will fail if all relevant reports are classified as 'software failures'. One means of avoiding this problem is to include free-text descriptions that provide additional details about the particular characteristics of each incident. As we shall see, these can be searched using specialist information retrieval techniques. These systems must recognise similar classes of failures in spite of the different synonyms, euphemisms and colloquialisms that are provided in free-text accounts of 'bugs', 'crashes', 'exceptions' and 'run-time failures'. In general, however, users may be forced to manually comb through each recorded software failure to extract those that relate to floating point exceptions.

The US National Co-ordinating Council for Medication Error Reporting and Prevention's Taxonomy of Medication Errors provides an example of a more fine-grained approach to incident classification [581]. This contains approximately 400 different terms that record various aspects of adverse incidents. If this were embodied within a storage and retrieval system then it would enable analysts to pose a number of extremely detailed questions about both the causes and outcomes of adverse medication incidents. Such a sophisticated approach also implies a high-level of training for those who must complete any analysis. Systems that are based on a detailed taxonomy increase the potential for confusion and ultimately low recall because different classifiers may exhibit subtle differences in the ways in which they distinguish between the terms in the taxonomy. In consequence, a number of these systems exploit flow-charting and similar techniques to help analysts identify which fields relate to a particular incident. Figure 11.9 provided an example of this approach by illustrating the the Eindhoven Classification Model. Analysts must first decide whether the causes of an incident are primarily technical, organisational or 'human'. Each of these high-level categories is successively broken down into increasingly more detailed causal factors until the terminal nodes represent the ultimate classification that will be applied to an incident. A recent study of trained staff classifying incidents according to the MEDWATCH codes, described in previous sections, identified a vast array of potential pitfalls. These resulted in a recommendation that either the FDA consider funding the centralised coding of all event reports or that the coding scheme be redesigned to benefit from top-down decomposition techniques similar to those exploited by the Eindhoven approach. MEDWATCH codes could be merged with the coding systems and hierarchical structures available within the National Library of Medicines Unified Medical Language System with an additional hierarchical coding system for device problem coding [262]. Unfortunately, further problems complicate the use of these hierarchical coding schemes. As we have seen, many incidents stem from complex combinations of many different causal factors. It can be difficult to ensure that independent analysts will arrive at the same classification patterns even when they have access to such tools.

The problems of inconsistency in detailed classification schemes can be seen as a slightly esoteric concern. Many large-scale systems face the more prosaic problems of ensuring that staff provide all of the information that is required about an incident. These problems can be exacerbated when staff must search through lists of valid codes to ensure a correct classification for each data field. The US Food and Drug Administration expressed their concern about this issue in their User Facility Reporting Bulletin. This provides feedback to the individuals and organisations who contribute information about device related failures. In an article entitled 'THOSE CODES' they describe how:

> "The final Medical Device Reporting regulation became effective July 31, 1996. Since then, Food and Drug Administration (FDA) staff have observed numerous errors and omissions in the MDR reports submitted by user facilities to report device-related deaths and serious injuries. These errors cause major gaps in FDAs adverse event reporting database, and may also delay manufacturers failure analyses while the manufacturers contact user facilities for additional information. FDA plans to send letters to those user facilities that have submitted incomplete mandatory forms (3500A) to request they file supplemental reports". [275]

One obvious means of addressing this problem is to ensure that analysts receive explicit training in the application of a classification or coding scheme. The APSF operates what is arguably the most elaborate of these systems [36]. Their training scheme is designed to provide experience of each of the 95 different options that can be coded into the AIMS system. The existing database is used to ensure that trainees meet a representative cross-section of scenarios as they learn how to use the classification scheme. They also work under the supervision of an APSF 'accredited coder' even though most of the course is conducted remotely using email and telephone contact. The training consists of an initial orientation session that covers the general motivation behind incident reporting and monitoring. They are taught how to install the associated AIMS+ software. They are then guided through an initial data entry and classification exercise. There then follow three different levels of training. In the first level, the trainee must code a sample collection of fifty incident reports. During this process they can request as much help as they consider to be necessary from the accredited tutor. The results of this classification exercise are then reviewed and graded by their supervisor who will provide appropriate feedback to the trainee

Second level training involves the coding of another fifty incident reports. In contrast to the previous exercises, however, the trainee is encouraged not to seek help from their supervisor unless absolutely necessary. These are again reviewed and graded before feedback is provided. In order to progress to the final level of accreditation, they must achieve a 60% 'pass rate'. This requirement is highly significant. It reflects the recognition that it will not be possible to expect or achieve 100% agreement between different analysts immediately after a period of relatively intensive training in a particular classification scheme. If the trainee fails to satisfy this 60% requirement then they must repeat the second level training with a further bach of fifty incident reports. Trainees are permitted three attempts at this second level before the tutor is required to refer the candidate back to their sponsoring organisation. The third level of training involves a final set of fifty sample incidents. Individual discrepancies in the coding are reviewed by the supervisor and the trainee. The expected pass rate is now raised to 75% before the trainee can graduate from the course [37].

Such training undoubtedly provides important support for the codification of incident data. Unfortunately, it can be both time consuming and expensive. These factors act as powerful disincentives for many organisations. There are also doubts about the long-term effectiveness of such training. Even in Sentinel systems, where additional resources are targeted on a few 'case study' organisations, it can be difficult to demonstrate the success of such initiatives. The FDA identified problems including "lack of coding (estimated at 50% of incoming reports), incorrect coding, and use of codes that are too general to be useful (e.g., device malfunction)" [262]. The success rate for organisations who were outside of this select group can be expected to be correspondingly lower.

Chapter 11 identified problems that affect the use of coding schemes to inform the causal analysis of adverse occurrences and near-miss incidents. Causal factors change over time as new systems and working practices are introduced. For instance, the introduction of microprocessor controlled infusion devices has created the potential for incidents that could not have happened in the past [182]. Similarly, classification schemes may also change as new ideas are developed about the underlying problems that lead to human 'error' and system 'failure'. For instance, our understanding of the impact of workload on human decision making has changed radically over the last decade [426]. The application and development of a reporting system can also help to identify improvements to existing classification systems. Many coding systems provide analysts with the opportunity to state whether or not they believe that any necessary information has been omitted from their classification. The feedback received from these submissions can be used to distinguish data that cannot be included within a classification scheme from information that was simply overlooked during the coding process. For example, the APSF training scheme, mentioned above, was structured around a Generic Occurrence Classification system (GOC). In 2000, this was was updated to GOC+. The introduction of GOC+ was supported by a computer-based classification system. The interface to this tool leads the analyst through a process that is intended to collect all of the relevant information that is required for each type of event. The changes were intended to increase the scope and content of the system:

> "Since incident monitoring began, the APSF has learned a great deal more about the factors involved in incidents in healthcare. In order for the GOC to remain a relevant classification tool, this additional knowledge has been incorporated into the classification. Another priority in the development process was to improve coding consistency, accuracy and timeliness. By analysing the issues that influence consistency, accuracy and timeliness, the development team was able to focus the development on managing these issues." [34]

Changes to any classification scheme can create considerable problems for the maintenance of an incident reporting system, In particular, it can be difficult to ensure that all data is indexed in a consistent manner. For instance, they may already be incidents in the database that provide examples of new classification concepts. In such circumstances, analysts may be forced to manually reclassify thousands or hundreds of thousands of existing records. This is often impossible. In consequence, many incident collections become partitioned by the coding schemes that were used to compile them. Separate queries may have to be performed for records that were gathered before and after an update. Statistics of the form X% of all incidents were caused by Y will have to be parameterised by the duration of the data-set that supports this analysis, even though there may be data that could confirm this analysis over a longer time period. A number of reporting systems have, therefore, attempted to develop computer-based tools that will guide analysts through the task of converting between coding formats. None of these systems can, however, entirely remove the need for manual intervention when new data is required by revisions to an existing classification system.

14.4.2 Lexical Information Retrieval

The previous paragraphs have summarised the problems of data maintenance that can arise when reporting systems rely upon relational systems. Further problems restrict the utility of these systems for end-users. In particular, it can be difficult to overcome the problems associated with query formation, both in terms of the knowing what to ask and how to ask it. These potential limitations can be addressed through the development of user interfaces that hide the underlying relational model:

> "The semantic query system in AIMS 2 release 2 will enable users to drill down into the data without having to understand the underlying database structure. We will also include some basic data mining facilities that allow users to contrast and compare rates across locations, incidents, etc." [39]

These approaches create additional problems. By hiding the underlying model in a relational system, it can be difficult or impossible for users to learn how to form their own queries. In consequence, they are restricted to those questions that have been anticipated and are supported by the systems developers.

Information retrieval tools provide powerful alternative mechanisms for searching large collections of unstructured data. Brevity prevents a complete exposition of the many different techniques that have been exploited by these systems, for a more complete review see Belew [71]. In contrast, the remainder of this section focuses on lexical information retrieval techniques. This decision is justified by the observation that these techniques have had the most widespread impact on commercial retrieval systems. The following section examines case based reasoning approaches that provide a point of comparison with these more widespread tools. Lexical information retrieval systems, typically, rely upon a three stage process. Firstly, collections of documents are indexed. This process associates one or more keywords with a document. By automating this process it is possible to reclassify large collections as new incident reports are received by the system [791]. This is a significant issue, as we have seen the Food and Drugs Administration's MAUDE system receives between 80,000 and 85,000 reports per year. Automatic classification not only offers the possibility of reducing inter-rater reliability concerns but also can reduce the costs associated with manually classifying each incident within a relational database.

The second stage of the information retrieval process involves processing the user's query or information request. The intention is to identify terms that might be matched against the keywords that were identified for each document. This create problems. There can be a mis-match between the terms that a user exploits when looking for an incident and the keywords used during the indexing phase. There is also a danger that a retrieval system will under-value those terms that the user perceives to be the most significant in their query. In such circumstances a request for 'software failures in surgical procedures' might focus on surgical incidents rather than the more detailed criteria for 'software failures'.

The final stage searches through a collection to identify matches between the terms in the users' query and the keywords that index each report. Unfortunately, users tend to form general queries that match many potential documents even when they have a relatively precise information need. In consequence, it is likely that an initial request may have to be iteratively refined as users search through large scale incident databases. This search process can be supported by relevance feedback techniques. The user indicates which of the proposed documents were actually relevant to their query. This information is then used by the retrieval system to improve subsequent searches. For instance, greater weight can be placed on any future matches between the terms in a query and the keywords of documents that the user has recognised as being relevant to a previous query involving those terms.

Information retrieval tools have supported numerous applications and are ubiquitous on the World Wide Web. It is, therefore, surprising that they have not been more widely adopted to support incident reporting systems. One explanation for this is that they cannot, in their pure form, be used to collate the statistics that are more easily extracted using relational systems. In a relational database, incident reports are classified according to the detailed components of a data model. It is possible to provide particular percentages for the numbers of incidents within each pre-defined category. In contrast, information retrieval systems avoid the pre-defined data models that have been criticised in previous paragraphs. Information retrieval tools make inferences based on the terms in a query and keywords associated with a document to determine whether or not it is relevant to the user [71] Many of these inferences are based on heuristic algorithms that cannot be guaranteed to satisfy the users' information need. Information retrieval systems make incorrect assumptions about the content of the document being retrieved and about the nature of the user's request. In consequence, it is difficult to rely upon the number of items retrieved by a query when generating statistics about the frequency of particular incidents. Further manual analysis must be performed to ensure that the retrieval tool has correctly identified all relevant incidents. As we shall see, this additional analysis can involve two different tasks. It is important to filter out any irrelevant 'hits' from the retrieved documents. This can have profound consequences for incident reporting systems. There may be insufficient resources to manually search through the many spurious matches that can be returned by some information retrieval tools. Conversely, it is important to ensure that any relevant documents have not been missed by the retrieval tool. This is a significant problem because users may fail to recognise a pattern of previous failures if similar incidents are not being detected by a retrieval system.

Information retrieval tools avoid the constraints of rigid data models by focusing on lexical features of documents. Relevant documents can be identified by looking for similarities between the words that are used in a query and those that are contained in an incident report. For instance, if the user issued a request to find 'all incidents of computer failure' then the retrieval system would look for any reports containing the words 'computer' and 'failure'. This example illustrates the potential strengths of this approach. Users can compose queries that do not require any understanding of an underlying relational algebra. This example also illustrates many of the problems that complicate this approach. Firstly, the retrieval system will have to strip out 'noise words' from both the query and the incident collection. This is important if any match is not to be overwhelmed by commonly occurring words such as 'and' or 'the' that occur in almost every sentence. Secondly, any implementation will be forced to process the query in order to recognise lexically related terms in the document set. The query term 'computer' should also match 'computers', 'computerised' as well as 'computational'.

Information retrieval, typically, depends upon the identification of concepts and terms that can be used to discriminate between the items in a collection. Words that commonly occur in all of the documents within a collection are unlikely to provide useful information about these concepts. For instance, 'function words' such as 'it', 'and', 'to', are necessary for the construction of grammatical sentences. They have a relatively high frequency because of their grammatical role but provide little help in identifying the content of a document. Other terms can be regarded as noise within particular systems. For instance, words such as 'clinical' or 'doctor' occur in most medical incident reports. If they were used as document keywords then a significant amount of indexing space and retrieval time would be spent filtering through values that are unlikely to help users discriminate between large-scale collections of incident reports. Unfortunately, we cannot simply strip out 'noise words' based on their frequency alone. For instance, the FDA's MAUDE system yielded more than 3,000 matches for 'software' incidents in January 2002. Such terms cannot be regarded as 'noise' even though they appear in many documents. They provide critical information about the nature of the events that they describe. Many information retrieval systems, therefore, rely upon a *negative dictionary* rather than raw frequencies [71]. These enumerate the words that can be ignored during the retrieval process. These include standard lists of function words. Negative dictionaries can also be supplemented by domain dependent lists provided by the end users of the system. For instance, a variant of the MAUDE system might deliberately exclude 'clinical' and 'doctor' as potential keywords. Clearly, the content of negative dictionaries can have a profound impact upon the performance of an information retrieval system. They must, therefore, be validated in consultation with the end-users of the system. The content of such dictionaries must also be reviewed as the nature of incidents, and hence of the language that is used to describe them, will change over time.

It is important to identify common concept in queries and documents even though they may not contain exactly the same lexical forms. One means of achieving this is through the use of stemming algorithms. For example, a query might contain the word 'error' and an incident might contain 'errors'. Any indexing must be robust to such plural forms. It must also consider variants, for example by deriving 'woman' from 'women'. Fortunately, there are standard techniques, including Porter's stemmer, that can be used to address these potential problems [71]. They extend the ability of the search engine to identify potential matches between lexical terms. They also reduce the number of keywords that are associated with documents. Plurals are stripped out, only the singular 'roots' are retained.

As mentioned, many information retrieval systems exploit the notion of 'inverse document frequency' as a means of identifying useful keywords. Rare words provide better discriminators than more frequent terms. In consequence, many retrieval systems will revise the weightings associated with particular keywords whenever new documents are entered into the system. Changes in the pattern of language used to describe incidents or in the underlying causes of adverse events can be reflected by changes in the weightings associated with particular terms. This creates a paradox in which the increasing frequency of particular incidents might result in lower weightings within the retrieval system. There are further complications. If we consider a document containing the term 'software failure' then this might provide a useful index within a collection of incident reports about medical adverse events. It would not, however, provide useful information in a collection that was entirely devoted to medical software failures. From this it follows that the discriminatory value of any index is determined by its ability to distinguish between the contents of that document and the other items in a collection. An extension of this argument is that the importance of any keyword for a document is determined not by the absolute frequency of that keyword within a collection but by the relative frequency of that keyword within the document compared to the frequency of the term throughout the collection as a whole. A word that occurs frequently within a collection can still provide valuable information about a particular document if it occurs even *more* frequently in that report. In other words, the background or 'noise' frequency of a work can be used to identify a threshold value. This can be distinguished from the signal value of the word if its frequency exceeds this limit [71].

The previous approaches focus on individual keywords. In contrast, a number of retrieval systems rely upon vectors of terms both to characterise queries and to index items in a document collection. In this view, each keyword represents a different dimension along which to compare a document to a query. The following vectors illustrate this technique using binary values. A keyword is either present or absent from a document. Variants of this approach rely upon weightings to indicate how often a particular word appears or how 'significant' that word might be in determining relevance within a collection of incident reports:

	Keyword 1	*Keyword* 2	*Keyword* 3	...	*Keyword* *N*
Document 1	0	0	1	...	1
Document 2	1	0	1	...	0
Document 3	0	1	0	...	1
...	1	0	1	...	1
Document *N*	0	1	0	...	1
Query	1	0	1	...	0

The simplest approach would be to take the inner product of the query and document vectors as a metric of similarity. However, vector-based information retrieval systems can go beyond the isolated use of keywords to look for patterns in a document collection. Matches may be based not simply on the direct relationship between a query and the document that matches it best but also on the transitive relationship between that document and other similar reports. The components of a query can be expanded to include keywords that are not explicitly mentioned by the user but which are also common to those documents that best match the users query. For example, a user might issue a request to identify incidents involving 'catheter' and 'lines'. It might be observed that many other reports which contain the word 'catheter' do not contain the word 'tubing' but do contain terms such as 'tubing'. These partial matches might therefore be offered to the user during subsequent interaction with the system.

Unfortunately, vector-based approaches also suffer from a number of problems. Most users queries yield very few keywords and so their vectors can potential match a large number of documents in the collection. One solution to this is to construct a query vector both from the users most immediate request and from the keywords that have been extracted from previous search tasks. This creates the problem that documents which were incorrectly returned during previous sessions will continue to be returned during future interaction. Vector based techniques must also account for the problem of document length normalisation. This arises because longer documents are more likely to contain more keywords that shorter ones. Hence there is a greater likelihood that they will be returned in response to most queries. In incident reporting systems this relates to a tension between scope and verbosity. It can be difficult to distinguish between lengthy accounts that describe a large number of complex failures and those that simply use 'more words'. The Swiss Anaesthesia Incident Reporting System (CIRS) provides good examples of this tension where reports of similar incidents involve IV lines range from under 100 words to over 1000 [755]. There are a variety of potential solutions to the problems of document length normalisation. In the CIRS system, this is less of an issue given the relatively small number of additions each month. In larger-scale reporting schemes, lengthy documents can be divided and indexed separately to ensure that keyword vectors reflect the changing content of each section. If the resulting vectors are very similar then they may be merged to prevent the generation of unnecessary indices. In practice, however, this leads to a host of further problems [71]. Brevity prevents a full analysis of techniques for document length normalisation and the interested reader is directed to Singhal et al [743].

The previous paragraphs have introduced some of the main issues that arise during the development of information retrieval systems. It should be apparent that this term covers a broad range of different approaches. Many of these techniques have still to be applied to support the development of incident reporting systems. There have, however, been a number of recent attempts to extend the benefits that these systems provide to support search and retrieval tasks amongst large collections of occurrence reports. In particular, the FDA have pioneered the use of some of these approaches in the web-based interface to their MAUDE reporting system for incidents involving medical devices [272]. As we have seen, this system is based around a relational database using techniques that were described in previous section. It also provides access through the Verity free-text search engine. This relies upon a lexical analysis that has much in common with the information retrieval techniques described in previous sections. From the users' perspective, they can issue restricted free-text queries rather than being forced to compose more complex sentences using SQL syntax. Initially, users of the Verity interface to MAUDE are encouraged to enter either a single word, such as Catheter. This will yield only those incident reports that contain the exact spelling of the word that is entered. Alternatively, users can enter an exact phrase, such as Catheter line. This will yield records in which those words appear in the exact order specified by the query. Users can also perform searches involving multiple words connected by the AND operator, such as Catheter AND tubing. This retrieves records that contain both search words in any order and any location in the text being searched. The initial Verity interface provides users with information about how to perform these relatively straightforward lexical queries. The FDA also provide guidance on how to perform more complex retrievals. The OR operator can be used to find reports that contain one of two search terms. For instance, pregnancy OR folate returns documents that contain the word pregnancy or folate but not necessarily both. Parentheses can be used to form more complex queries. Quotation marks can also be used to explicitly denote that a literal match should be performed. Users can select documents that contain both pharmaceutical companies and stock by entering AND ('pharmaceutical companies', 'stock'). The , comma operator returns documents containing at least one of the words specified using a ranking approach. The FDA's implementation returns an ordered list; incident reports that contain the most occurrences of the keywords are given the highest rank. There is, however, no attempt to exploit the length normalisation algorithms mentioned in previous paragraphs.

It is also possible to create queries using the NOT operator. Ideally one might like to pose queries of the form NOT software to return every non-software related incident. The unrestricted use of such queries would create considerable computational overheads. This undermines variants of the indexing strategies described in previous paragraphs and would result in a form of exhaustive search over several hundred thousand records. Verity will, however, attempt to execute unrestricted queries involving the NOT operator. Issuing a request for NOT software in January 2002 returned more than 7,000 MAUDE records before the system ran out of resources and stopped the request. In contrast, the FDA recommend that users form queries that restrict the negated search term. The NOT operator 'finds documents containing the word that precedes it but that do not contain the word(s) that follows it'. For instance, pregnancy NOT folate yields incident reports with the word pregnancy but excludes any document that also contains folate.

The NOT operator demonstrates that many free text search facilities are not as 'intuitive' as they might first appear. They do, however, support the notion of proportionate effort. It is possible to perform literal keyword searches with minimal assistance. More complex query formation involves some additional thought. It might be argued that the implementation problems surrounding negated queries demonstrate that lexical forms of information retrieval offer few benefits beyond those provided by relational databases. This argument can, however, be challenged. In the case of relational databases, users must consider both the semantics or a range of relatively complex operators and the underlying data model that will be different for each database. In the case of lexical information retrieval tools, the user only has to understand the underlying concepts associated with particular operators. The proponents of these systems also argue that, in contrast to relational databases, most of the key concepts can be formed inductively without explicit training. Over time users will learn about the efficiency problems associated with unrestricted negation as they experience significant delays in processing their queries.

Previous paragraphs have described how the Verity retrieval tool searches for literal matches with the terms used in a query. This can create significant problems for many users. In particular, it can be difficult to search for all incidents involving particular manufacturers. The FDA acknowledge that 'when searching on company names, the search does not include variations of spelling or use of symbols such as hyphens, slashes, etc' [268]. However, the problems associated with exact literal match algorithms are exacerbated by the difficulty of data validation in large scale incident reporting systems. During the preparation of this book, I found numerous instances in which the names of manufacturers had been misspelt in the sections of the incident report that were searched by the Verity system. In consequence, users must exploit literal search facilities to identify incident reports that contain the correct spelling for a device manufacturers. They must then form additional queries to check whether any reports have been missed because those names were mis-spelt. This problem can be avoided in relational systems where manufacturers must be associated with one of a number of pre-defined attributes. The Verity system does, however, provide additional operators that can be used to address this limitation of using free text data during the analysis of incident reports. The ? question mark provides a wild-card that can represents any single character. For instance, the query ?ietermans would locate documents containing the words Viertermans, Fiertermans, Giertermans and so on. In contrast, the * asterisk represents one or more characters. A query of the form corp* would return documents containing corporate, corporation, corporal and corpulent.

The Verity interface to MAUDE also provides users with access to some of the stemming facilities that have been described in previous paragraphs. Queries that exploit this facility must include key terms using single quotation marks. For example, the query cath' finds catheter, cathlab, cathode and cathodic among others. This explicit approach to query formation using stemming can be combined with the <MANY> operator to count word densities in FDA incident reports. For instance, the query <MANY> cath' produces a ranked list in which the first document contains the most occurrences of words with the cath stem. In contrast to the comma operator introduced in previous paragraphs, Verity's <MANY> queries do perform length normalisation. Hence the FDA advise that 'a longer document that contains more occurrences of a word may score lower than a shorter document that contains fewer occurrences'. Verity offers a range of more complex operators that can be used to search for words within particular sections of an incident report. For example, the < NEAR/N> operator can be used to find documents that contain words within a specified distance of each other. For example, the query, balloon < NEAR/10> rupture, would locate all documents with the terms balloon and rupture within ten words of each other. Similarly, < SENTENCE> and < PARAGRAPH>

will find documents in which the specified terms are in the same sentence or paragraph.

The previous paragraphs have focussed on the facilities that the FDA's Verity tool provides for lexical information retrieval across the MAUDE incident collection. These facilities are built upon partial or literal matches between keywords in a document and the terms in a query. There are, however, a range of information retrieval techniques that make inferences about potential matches that go beyond the keywords that appear in a document. Many of these approaches rely upon thesauri that represent the relationships between keywords. In consequence, if there are few literal matches between a query and the documents in an incident collection then retrieval tools can look for matches between a query and other keywords that are in some way related to those in the document. Alternatively, the users' query can provide the basis for additional searches using terms related to those in the original request. Thesauri have been extended to include the following relationships:

- synonymy. Two expressions are synonymous if the substitution of one for the other does not change the interpretation of a sentence. For instance, cardiopulmonary resuscitation is synonymous with artificial respiration and heart massage. This relationship can also be used to connect acronyms to their associated terms. Hence CPR is related to cardiopulmonary resuscitation. Synonymy can also be used to capture conventional or authoritative keyword that replace less favoured terms used within a document or query. For instance, amyotrophic lateral aclerosis or ALS might be preferred to Lou Gehrig's Disease. Such relationships are critical in natural language queries that search for similar incidents describe from different perspectives. A variety of terms can be used to describe the same concepts depending on the geographical location of an incident or the functional role of the reporter.

- antonymy. This relationship is less commonly supported than the other forms in this list. Antonymy represents a pair of words which are related by an associative bond. These associations are often validated in empirical studies, or word associated tests, involving the potential end-users for a retrieval system. Antonyms are often revealed to have an opposite semantic relationship to the probe terms used in th studies. Hence many people will respond with the term victory when promoted with the word defeat and vice versa. As we shall see, there have been few attempts to apply this form of relationship to support information retrieval within an incident reporting system. It can, however, be argued that such techniques might be used to identify successful instances of a procedure rather than previous failures.

- hyperny/hypony. A hypernym designates a class of specific instances. Y is a hypernym of X if X is a (kind of) Y. In contrast, a hyponym describes a member of a class. A hyponym inherits all of the features of the more general hypernym and adds at least one feature to distinguish it from the high-level concept. For example, Lymphoma can be treated as a type of Neoplasm [71].

- meronymy/holony. Meronyms are constituent parts or members of something else. Hence, X is a meronym of Y if X is a part of Y In contrast, a holonym is the whole of which the meronym names a part. For instance, the cecum is the first part of the large intestine and is hence a meronym for the larger structure. Any query about incidents involving procedures on the large intestine might also return procedures involving the cecum.

This partial list only provides an indication of the relationships that can be used to expand on the keywords derived from a query or used to index a document. These techniques have not, however, been widely applied in either incident or accident reporting systems. Carthy has recently begun the first systematic examination of thesaurus-based retrieval techniques for incident reports in a project funded by the Irish Government. His work builds on automated topic detection and tracking systems. These applications enable their users to identify common threads amongst the publications of news and broadcast media. It also exploits the development of domain specific taxonomies, such as National Library of Medicine's Medical Subject Headings Thesaurus and the National Co-ordinating Council for Medication Error Reporting and Prevention's Taxonomy of Medication Errors [581]. Although these systems have been developed to support other applications, they can be directly applied to support the retrieval of incident reports [152]. It is also hoped that Carthy's work will encourage the development of techniques that are specifically intended to detect patterns of failure in reporting schemes. In anticipation of this research, the FDA's Verity interface provides access to more general facilities. The <THESAURUS> operator expands a search based on synonyms of the word(s) in a query. The example provided in the FDA documentation is that the query <THESAURUS> altitude will yield documents that include the terms height, elevation and altitude. The <SOUNDEX> operator expands the search to include words that 'sound like' the term(s) in a query.

The previous analysis of the FDA system again illustrates two key issues. Firstly, that advanced information retrieval systems can make powerful use of thesauri and similar techniques to make inferences about relevance that go well beyond the terms contained in either a query or a particular document. Secondly, that some training may be required if users are to fully direct or control the facilities that these techniques provide. In particular, the use of either the <SOUNDEX> or <THESAURUS> operators can lead to a rapid rise in the number of hits that are detected by the system:

> "If your full text search is broad, you may be attempting to retrieve more then the system limitation. If this happens, you will receive a message indicating that your record retrieval is incomplete. The system is not capable of retrieving any missing records over the limit". [268]

The following paragraphs will describe ways of measuring the adverse consequences of such information 'overload'. It is, however, important to consider recent attempts to control this problem by integrating information retrieval techniques and relational databases. For instance, Chapter 5 has already described how many incident reporting forms combine check-box questions with more open-ended questions that can be answered using free-text. Computer-based systems can, therefore, be developed to implement the strongly typed check-box information using relational techniques; each check-box represents a value for one of the attributes in a relation. The same system might combine this with lexical techniques for information retrieval so that users can search in a less directed fashion over the free-text descriptions of adverse events and near-miss occurrences. In theory this combine approach can offer numerous benefits. For instance, statistical returns that require deterministic answers to particular focussed queries can still be conducted over the data that is stored using a relational database. Less directed 'information mining' operations can exploit the free-text areas of each report. This complementary approach can also help to control the information 'overload' problem. If a thesaurus is used to expand query or document keywords then the mass of potential returns can be filtered by restricting the search to incident reports that match particular relational attributes. For example, the use might direct a request to find all incident reports relating to <THESAURUS> CATHETER so that it was only evaluated over reports filed by device manufacturers or, alternatively, by end-user facilities. If the report contributor were recorded as a field in the relational component of the system then users might be relatively confident that their query was restricted in the desired manner. Such a filtering would be less easily achieved using lexical variants such as NOT END-USER because the free-text accounts might not all have used the term END-USER and thesauri-based techniques do not guarantee to find all possible forms of synonym that may have been used by every contributor.

The FDA's MAUDE system provides a partial integration of their relational model and the Verity retrieval system. The lexical analysis is restricted to free-text areas of the device reports and does not cover any fields that are 'encoded' using numeric or other identifiers. The database elements that are examined by the Verity facility include: MDR Report Key; Manufacturer Name; Distributor Name; Brand Name; Generic Name; Model Number; Catalogue Number; Product Code and Adverse Event or Product Problem Description At first sight this might appear to offer the form of integration mentioned in the previous paragraphs. It is possible to use Verity to search over the attributes of relations within the MAUDE database. The FDA note that 'the Full Text Search cannot be combined with any other search options on the MAUDE search page' [268]. It is not possible to apply Verity to a subset of incident reports that have been filtered using the query language provided by the relational system. Hence it can be argued that the FDA provide a form of data-level integration rather than a full system-level integration. Information can be shared between

Verity and the relational format but both systems cannot easily be used to construct hybrid queries.

This section has argued that the information 'overload' problem might be overcome by using relational queries to filter the incident reports that are examined using lexical approaches to information retrieval. Unfortunately, this can only be a partial solution to what is a more complex problem than has previously been suggested. In large scale systems, relational filtering may still yield enormous numbers of incident reports in response to thesaurus-based free text queries. It is important to recall that MAUDE includes almost 400,000 records at the start of 2002. Analysts would still have to invest considerable time and energy to identify common features even if a query returned only 1% of the reports in the system. One solution to this problem would be to develop more precise data models within a relational system so that users could filter on more detailed features of an incident. This is unsatisfactory because increased discrimination tends to be achieved at the price of increased complexity. Alternatively, lexical analysis can be focussed more tightly to filter out spurious matches. For instance, by restricting the use of a thesauri it is possible to focus on a narrow selection of synonyms. Unfortunately, this increased precision will also typically result in worsening recall rates. The lexical analysis will miss reports that contain related terms and concepts, which were excluded by the narrow associations provided in the thesaurus.

Precision and Recall

Precision and recall are concepts that are used to assess the performance of all information retrieval systems. In broad terms, the precision of a query is measured by the proportion of all documents that were returned which the user considered to be relevant to their request to the total number of documents that were returned. In contrast, the recall of a query is given by the proportion of all relevant documents that were returned to the total number of relevant documents in the collection [220]. It, therefore, follows that some systems can obtain high recall values but relatively low precision. In this scenario, large numbers of relevant documents will be retrieved together with large numbers of irrelevant documents. This creates problems because the user must then filter these irrelevant hits from the documents that were returned by their initial request. Conversely, other systems provide high precision but poor recall. In this situation, only relevant documents will be returned but many other potential targets will not be retrieved for the user.

Belew [71] defines precision and recall in terms of the intersection between two sets:

$$Recall \equiv \frac{\mid Retrieved\_documents \cap Relevant\_Documents \mid}{\mid Relevant\_Documents \mid} \tag{14.20}$$

$$Precision \equiv \frac{\mid Retrieved\_documents \cap Relevant\_Documents \mid}{\mid Retrieved\_Documents \mid} \tag{14.21}$$

This is illustrated in Figure 14.5, which provides a high-level sketch of the relationship between precision and recall in information retrieval systems. Image a) reflects a query that achieved both high precision and high recall. Most relevant documents and no irrelevant documents were retrieved. In contrast, image b) represents high precision but poor recall. Only relevant documents were returned but many potential 'hits' were missed by the system. Image c) shows poor precision and high recall. Many irrelevant documents were retrived and hence the system is imprecise. In contrast, the query yielded all of the relevant documents so the system showed good recall. Finally, image d) shows both poor precision and poor recall. The query yields many irrelevant documents and retrieves very few that provide the required information.

Although the concepts of precision and recall are widely used in the evaluation of information retrieval systems there remains considerable disagreement about how to measure them in practice. These measures do not simply relate to system performance, they also relate to the corpus or collection that is being used. A system that achieves good recall rates on one set of documents may not achieve the same level of performance on another. This is particularly true for systems that rely upon thesauri. The meaning of key terms may differ considerably between domains and hence the system will have to be tailored to reflect differences in usage. For example, in the time series analysis of cardiovascular data the term 'leakage' is used to describe a loss of power from a

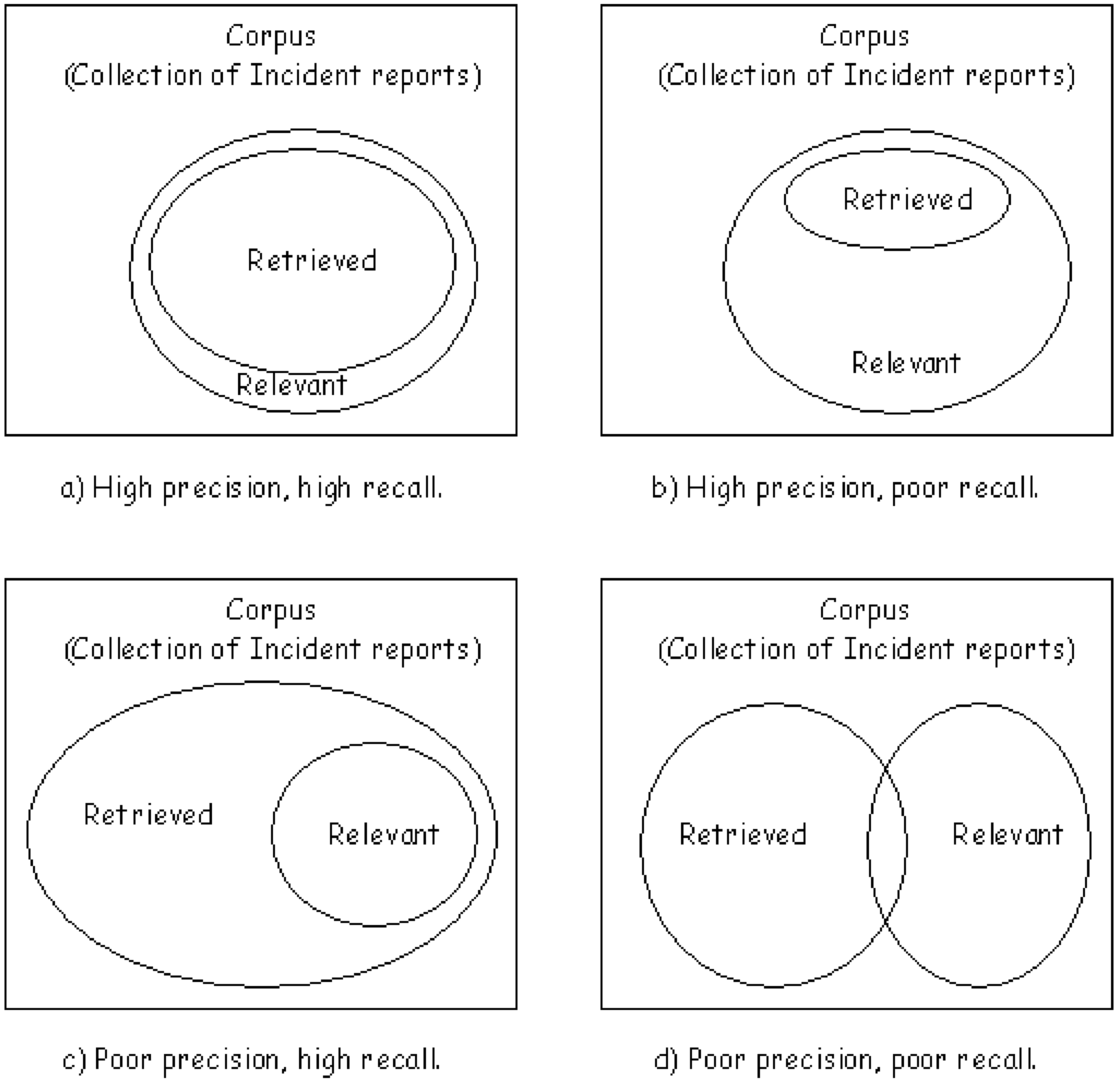
Corpus
(Collection of Incident reports)
Retrieved
Relevant
Corpus
(Collection of Incident reports)
Retrieved
Relevant
Corpus
(Collection of Incident reports)
Retrieved
Relevant
Corpus
(Collection of Incident reports)
Retrieved
Relevant

a) High precision, high recall.

b) High precision, poor recall.

c) Poor precision, high recall.

d) Poor precision, poor recall.

Figure 14.5: Precision and Recall

frequency band to several adjacent spectral lines which is typically due to the finite data set over which the periodogram is estimated. This is very different from more general applications of the term and hence appropriate relationships between synonyms would have to be explicitly encoded into an information retrieval system. Even if these relationships were encoded in a way that ensure good performance in the medical domain, there is no guarantee that the same system could easily be ported to, for instance, aviation. More research is urgently required to determine whether the linguistic characteristics of incident reports within these different fields can be used to support many of the retrieval techniques, mentioned above. Even within a topic, performance can vary depending on the nature of the documents that are contained within a collection. It is for this reason that information retrieval tools are, typically, evaluated using standard collections that provide a 'gold standard' for performance comparisons. This creates some problems for safety managers who want to exploit this technology. It is far from certain that the recall and precision values that can be obtained from a 'standard' corpus in information retrieval research will be mirrored in the operation of an incident reporting system.

There are further complications. The previous description of Figure 14.5 depicted precision and recall as properties of a particular query. For example, image a) shows a query that results in both high precision and high recall. It is important to recognise, however, that recall and precision vary dramatically depending on the query that is evaluated. For instance, if a thesaurus based system recognised a number of synonyms for a keyword then it is likely to provide high recall values for any query involving that term. In contrast, poor recall rates might be anticipated if the same system were presented with a query that did not contain any recognisable keywords. In consequence, comparisons between the precision and recall rates for particular systems must often be made in terms of specific queries on a particular data-set. If this were not the case then misleading values might be presented for carefully chosen requests. This raises the very practical concern that safety managers identify a 'realistic' test suite when attempting to evaluate the relative merits of these search engines. They must also identify an appropriate set of queries that reflect the likely information requirements for the end-users of the system. As we have seen, these can be difficult or impossible to predetermine given that the nature of incidents will change over time.

There are further complications. The images in Figure 14.5 assume that it is possible to unambiguously determine whether documents are either relevant or irrelevant to a particular query. There is no 'fuzziness' in the membership of *Relevant_Documents*. This reflects a strong assumption within the information retrieval research communittee that cannot easily be maintained for most 'real world' systems [422]. In particular, it does not characterise search tasks involving collections of incident reports. In many cases, it is difficult to be sure whether or not a particular document is relevant to a particular query. To illustrate this point, Jeffcott has recently conducted a study in which Risk Managers in Scottish hospitals were asked to read 8 reports of medical adverse incidents ranging from a problem in the use of a Doppler Fetal Heart-Rate monitor through to a morphine overdose [395]. Each incident was selected by a consultant and a senior nurse to provide a broad cross-section of the incidents reported to their Unit. The Risk Managers were asked to associate each incident with a number of broad categories that might correspond to retrieval requests using a variant of the 'expressed preference sampling' procedure developed by Fischhoff, Slovic, Lichtenstein, Red and Combs [249]. In simple terms, they were asked to rank whether or not they agreed with particular statements about an event using a 7 point scale. The results of this study showed a marked reluctance to use the extremes of the scale. The Risk Managers were unwilling to state that particular incidents did or did not exhibit a strong relationship to the questions that were posed. These responses undermine the binary distinction between relevant and irrelevant documents that is often assumed to exist in validation techniques for information retrieval systems .

It is difficult to under-estimate the importance of precision and recall to the application of advanced search techniques within incident reporting systems. In most other areas, including web-based retrieval, the trade-off between precision and recall can be characterised as either a performance or usability issue. In incident reporting schemes, these characteristics have safety implications. Low-recall results in analysts failing to identify potentially similar incidents. This can lead to litigation in the aftermath of an accident. Failure to detect trend information in previous incident reports can be interpreted as negligence. Conversely, low-precision leaves analysts with an increasing manual burden. They must filter the irrelevant documents that have been identified as hits by the search engine. This will result in omissions and 'errors' if fatigue or negligence undermine the manual filtering.

In spite of the problems in assessing the performance of lexical information retrieval systems, it is likely that these applications will play an increasingly important role in the computer-based dissemination of incident reports. The reasons for this centre on the need to provide technological support for the sharing of incident information between and within heterogeneous organisations. As we have seen, lexical information retrieval systems support the notion of 'proportionate effort'. Simple queries can be formed in a relatively flexible manner with only a limited understanding of the underlying data representation. More complex queries can be formed providing users understand the basic mechanisms involved in lexical retrieval, such as the use of a thesaurus to identify synonyms for keywords. It is not, however, necessary for users to learn the specific data representation that are associated with different incident databases. This contrasts strongly with the use of relational systems where it is necessary to understand the underlying data model before users can construct well-formed SQL queries. Such learning overheads would be of limited importance if safety managers only had to access a single incident database. Over time, novice users will gain experience in using the relations that lie behind systems such as MAUDE. Unfortunately, the lack of standardisation within many industries has combined with the increasing availability of web-based information resources to create a situation in which safety managers may have to understand the underlying data models associated with several different reporting systems. For instance, the UK MDA uses a relational model that is quite different from that used to describe US medical incidents. A small number of international initiatives are beginning to address this problem. We have mentioned the GAIN programme within the aviation industry in previous chapters [308]. This is, however, focusing more on the analytical techniques and underlying technological infrastructure necessary to support information sharing. Limited progress has been made towards the development of integrated data models for incident reporting that might enable users to exchange and search information from competitor companies in a convenient manner.

Several further factors increase the likelihood that lexical information retrieval systems will provide the technological infrastructure to support the dissemination of incident data between different reporting systems. The commercial impact of the world wide web is arguably the most important of these factors. Rapidly identifying relevant documents amongst a mass of other data is a key business requirement for many of the organisations and individuals that use the world wide web. In consequence, many companies are investing heavily in the technologies that support these tasks. This has produced tools that enable users to perform interactive retrieval tasks involving many millions of documents. These commercial developments offer further benefits. It is important to recognise that most of the documents that are placed on the web are unstructured. They owe more in common to the natural language accounts that are amenable to lexical information retrieval than they do to the more rigid relations within a database model. If a user issues a request for information about a medical product, they cannot expect that every device manufacturer will format the pages about their products in exactly the same way. This analogous to the situation facing safety managers and regulators looking for patterns of failure across several incident reporting systems. They cannot assume that all of these systems will exploit the same relational model. In consequence, lexical information retrieval systems offer a flexible means of analysing incident reports produced in many different formats by many different agencies. As we have seen, however, these systems do not yield the deterministic results that are typically required by statistical analysis. The precise number and nature of incidents returned by any query will depend upon the thesaurus that is being used and upon the discriminatory value of keywords that will change over time [71]. In consequence, common interchange formats for relational databases still offer considerable benefits for the exchange of incident data. The development of such common formats or schemas will not, however, resolve the problems associated with inter-rater reliability in the assignment of particular values to the attributes in a relational model. In consequence, I would argue that lexical retrieval tools will continue to provide the only feasible means of creating multi-national incident databases within the near future. Many safety managers and regulators already use mass-market retrieval systems to search for mandatory occurrence reports that are routinely placed on the web sites of the

CAA, NTSB and similar organisations. As we have seen, however, these more general tools do not support the domain specific thesauri that can be used to extend the scope of particular searches to achieve improved recall and precision. Similarly, it can be difficult to ensure that these mass market tools only retrieve potential hits on recognised sites. A search on catheter and incidents returns advertising material from manufacturers, research advertisements from government organisations, general news items from publishers, collections of papers published by particular individuals and so on. In order to address these problems, we have developed a series of web-crawlers that restrict the keyword indexing of documents to incident and accident reports on named sites. The terms that are used in the indexing and retrieval process are based on interviews with safety managers and regulators within the domains that we are investigating, principally rail and aviation safety, and are tuned according to the lexical frequency of terms within the collections of incident reports that we are studying. For instance, queries involving the term 'CRM' will yield incidents mention 'Crew Resource Management', 'Communication Failure' and so on. The intention is that these systems will provide feasible means for safety-managers and regulators to search for patterns of failure across the pages of incident reports that are increasingly being published via the world wide web [285, 528].

14.4.3 Case Based Retrieval

The previous section has identified a number of limitations of relational databases and lexical retrieval systems. Relational systems, typically, use strictly defined data-models to structure the information that is recorded about an incident. The many different individuals who enter or retrieve data from these systems often only have a limited understanding of these models. Further problems arise when changes are made to the components of a relational model; it may be necessary to manually reclassify hundreds of thousands of existing records. Alternatively, lexical search engines can be used to identify related terms in many different incident reports. Stemming techniques and thesauri can be used to expand queries or documents so that retrieval does not depend on literal matches. These approaches also avoid some of the problems associated with relational data-models. Users can enter natural language descriptions of each incident. Requests can be expressed as (pseudo) natural language queries. Unfortunately, as we have seen, the matching processes depend upon the frequency of terms within a collection. It may also be affected by relatively small changes within a thesaurus. In consequence, lexical approaches cannot easily provide the types of statistical returns that are required by regulatory organisations. A number of further problems relate to the precision and recall provided by these retrieval techniques. Precision is defined as the proportion of documents that the user considers being relevant within the total number of incidents that are retrieved. Recall is defined as the proportion of relevant documents that are retrieved against the total number of relevant documents within the entire collection. Hence an information retrieval system may have high recall and poor precision if it returns a large number of the relevant incidents in the entire collection but these incidents are hidden by a mass of irrelevant incidents that are also retrieved. Another system can have good precision and poor recall if it returns very relevant incidents but only a small proportion of those that pertain to the topic of interest. Many users have great difficulty in composing free-text queries that achieve a desired level of precision or recall. Most searches provide a small number of appropriate documents with many more irrelevant references. This poor level of precision can be exacerbated by inadequate recall. It is rare that any single query will yield all of the possible references that might support a user's task. These limitations can be frustrating for the users of mass-market retrieval techniques, such as web-based search engines. They can have more profound consequences for incident reporting systems. There are clear safety implications if a search engine fails to return information about similar incidents. A pattern of previous failure may be hidden by the poor precision or inadequate recall of some retrieval tools.

Case-based reasoning techniques relax some of the strict classification requirements that characterise more traditional databases. They do not avoid the concerns over precision and recall that affect other information retrieval tools. However, they often provide explicit support for users who must issue queries to identify similar classes of incidents within a reporting system. In the past these systems have been used to support fault-finding in computer systems, the design of wastewater treatment systems and route planning for mail delivery [455]. Ram provides an overview of this approach; 'case-based reasoning programs deal with the issue of using past experiences or cases to understand, plan for, or learn from novel situations' [692]. Most of these systems are based around a four stage process. Firstly, problem descriptions are used to identify previous similar cases. Secondly, the results achieved by attempts to address these previous cases are passed to the user. Thirdly, some attempt is made to extrapolate from the results of previous cases to the likely outcome of a similar approach being applied to the current problem. Finally, a generalised representation of both the old and new solutions are entered into the system so that future problems might benefit from any insights obtained during the analysis of the current problem. Ram's general analysis of case based reasoning can be applied to illustrate some of the potential advantages that this technology might offer for the analysis and retrieval of incident reports. Each 'case' can be thought of as an incident report. The attempts to resolve those cases can be seen as the recommendations that were made following those previous incidents.

The central problem of case-based reasoning is how to generalise from the specifics of a new incident so that it is possible to recognise any underlying similarities with previous cases. This is not as straightforward as it might appear. It is possible to identify at least three possible outcomes for any search:

1. *Exact match.* Two incidents are identical. In particular, we might be interested in those incidents that share both common causes and consequences [456]. Such similarities should not be discounted as unlikely given the increasing scale of many reporting schemes.

2. *Local divergence.* We might also want to identify partial matches between a new incident and previous cases. Two incidents share the same causes but an additional event or circumstance during one of the incidents led to divergent consequences. Alternatively, two incidents might have the same outcome but different causes. This reflects the causal asymmetry noted by Hausman [313] and described at length in Chapter 11.

3. *Global divergence.* Two incidents have no apparent similarity. They stem from different causes and result in different outcomes.

Case-based reasoning exploits some of these distinctions. For instance, an exact matching offers considerable efficiency gains because two cases can effectively be treated as a single more general case during the final stage identified by Ram, described above. Local divergence can be used to generate new indices that distinguish between cases with, for example, different causes or consequences.

Cases can be represented in a number of ways. Keyword or feature vectors, introduced in the previous section, can be used to represent whether or not particular terms are relevant to an incident. For example, the following narrative describes an incident from the MAUDE collection. This individual case might be represented by a vector that indicates the presence of indicative terms such as 'software', 'upgrade' and 'package':

> "During in-house software testing (of an ultrasonic analysis package), the manufacturer discovered unexpected software behavior in the generic tool kit when waveforms were inserted, resulting in correct calculation for the tricuspid valve regurgitant orifice area measurement.
>
> The software problem was found during in-house software development. It occurs when the user attempts to make a specific calculation in the cardiac calculations package and is related to a formula error. The software error has been identified and was corrected in a subsequent revision of the system software. Actions taken include customer notification of problem and installation of software upgrade to affected systems. It is important to note that there was no reported adverse event to a patient as a result of this event."

Stereotypes can be used to identify patterns of failure between the individual incident vectors. Each stereotype can be represented by the terms that a domain expert or contributor might use to describe particular incidents. For example, a stereotypical report of a software failure might include terms such as 'bug', 'crash', 'program', 'error', 'upgrade' and so on. If an incident report contained these terms then the associated similarity measure would be incremented each time that they appeared. In the previous example, the software stereotype score would be incremented for the terms ‘program’, ‘error’ and ‘upgrade’ because these are mentioned in the MAUDE account. The case-based reasoner returns a ranked list of stereotypes based on these similarity metrics. It is important to stress that the ranked list might return high scores for more than one stereotype. This is appropriate given that a software failure might be compounded by operator error or another form of adverse event. Each stereotypes can also be associated with particular remedial actions. For example, if software failure was returned as the highest ranked stereotype then the user of the system could be prompted to consult a guidance document on recommended procedures for resolving such incidents. This illustrates how lexical retrieval techniques can be intergrated into a case-based reasoning system.

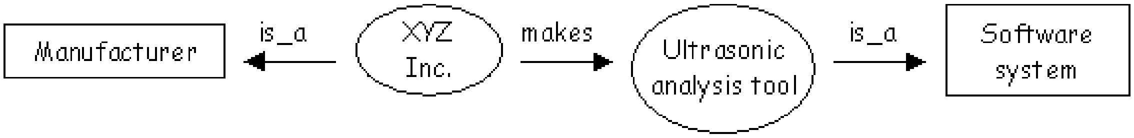

Figure 14.6: Components of a Semantic Network

The use of term-vectors is only one of several alternative approaches that can be used to represent and reason about common patterns in individual incidents. For instance, semantic networks can model an incident and more general aspects of the domain in which a failure occurs. In their simplest form, a semantic network can be thought of as a series of nodes and edges. The nodes represent objects and concepts in the domain of discourse and the edges represent relationships between them. For instance, Figure 14.6 represents part of the MAUDE incident report for the software failure that was cited in previous paragraphs. This diagram includes two different types of node. Rectangles are used to denote higher-level abstractions that may be common to many different cases. Software systems and manufacturers are likely to be involved in a more than one incident. In contrast, elipses represent particular instances of those abstractions. XYZ is a particular manufacturer, the company named in the report has been anonymised here. Similarly, an ultrasonic analysis tool is a particular type of software system.

Figure 14.7 extends the semantic net for the ultrasound software failure. As can be seen, the MAUDE narrative provides information about several different aspects of this incident. The failure mode was detected during a wave form insertion test. The problem was remedied by notifying the customers and by issuing a software upgrade. The fault might also have resulted in a patient injury. This diagram does not provide any information about a particular outcome for this incident. Figure 14.7 might, therefore, be revised to explicitly denote that nobody was injured as a result of this incident. This illustrates how the high-level abstractions in a semantic network can be used to provide an alternative to the lexical stereotypes, mentioned in previous paragraphs. Rather then relying on work frequencies to cluster similar incidents, semantic networks can be used to describe common relationships that characterise particular types of adverse event. Figure 14.8 shows how this can be done by removing all of the instance information from the previous semantic network. This leaves a high-level description not simply of the MAUDE incident that we have analysed but, more generally, of many different software-related failures.

Case-based reasoning systems can use the abstractions in Figure 14.8 in a number of different ways. They can be used like the components of a relational model to prompt uses for particular information whenever they enter information about an incident that seems to match with a particular stereotype. If, for example, the system determined that the new incident included a software-related failure then it might prompt the user to provide information about the detection method. This process can, in turn, contribute to the development of more appropriate abstractions. If a new incident was detected by an end-user facility then Figure 14.8 would have to be amended. The existing abstractions only consider Test Procedures as a means of detection. This process of case-based generalisation represents an instance of the final stage in Ram’s taxonomy of case-based learning, mentioned above [692]. It also illustrates how the development of semantic networks from

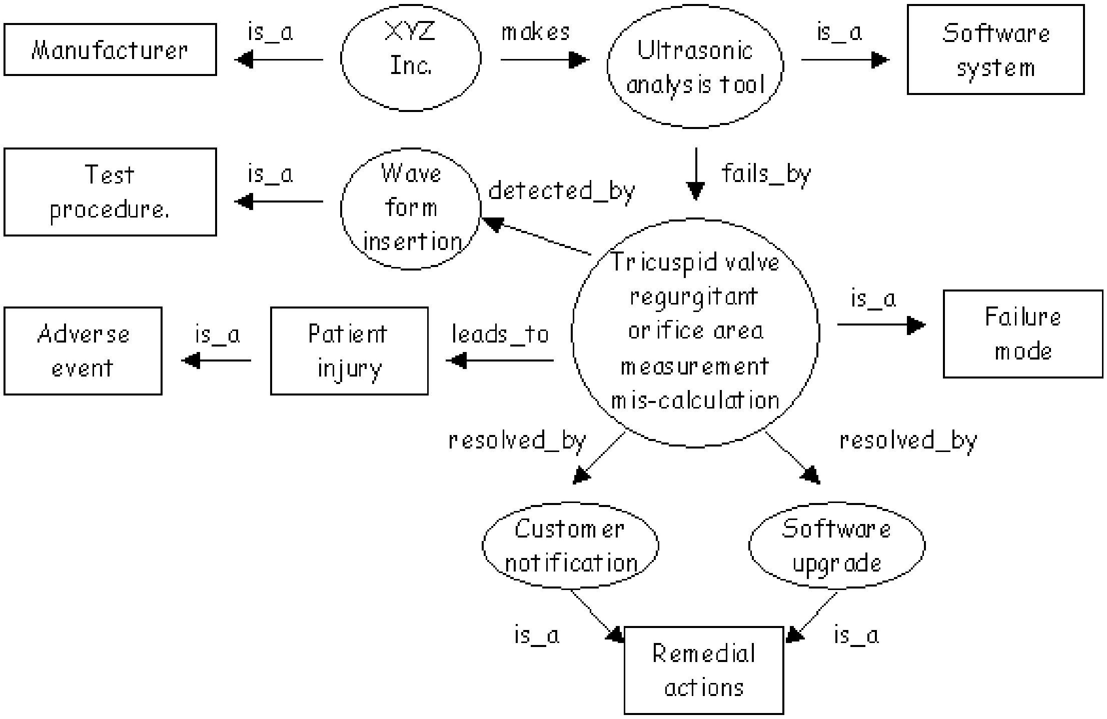

Figure 14.7: Semantic Network for an Example MAUDE Case

individual cases can help to create an ontology for particular types of incident. These ontologies provide a common reference point for the kinds of objects and relationships which characterise certain failures.

Figure 14.8 illustrates similarities between relational schemas and high-level abstractions from case-based reasoning systems. There are, however, strong differences between these two approaches. As we have seen, case-based learning systems are explicitly designed to cope with changes in the high-level models that represent previous failures. This contrasts with the costs that arise from changing the data model in a relational database. As we shall see, case-based reasoning systems also typically hide the detailed components of the underlying networks. Users are not expected to form complex queries that depend both on the underlying model and components of the relational algebra. Further differences stem from the matching algorithms that are used to determine whether a new case is similar to a previous incident. Both case based-reasoning systems and relational databases support instantiation or literal substitution. Similar incidents can be identified by looking for previous records with identical attributes. In addition, many case-based systems also exploit knowledge-based search techniques. These approaches extend the semantic networks shown in this chapter to support the thesauri-based approaches described in the sections on lexical retrieval techniques. For example, a search might be made through the previous cases to find incidents that were detected by tests which are synonyms of wave form insertion, such as wave form addition or wave form introduction.

There are many more complex variations on the general approach described in previous paragraphs. For instance, Kolodner pioneered many of the initial case-based reasoning techniques using a Dynamic Memory Model that was based on 'generalised episodes' [454]. These episodes form a hierarchical structure. For instance, at the highest level the MAUDE system describes episodes that relate to the failure of medical devices. These can be further sub-divided into episodes that describe software failures, human error and so on. Each 'generalised episode' is described in terms of norms, cases and indices. Norms are common to all of the cases indexed under a generalised episode. For instance, a normal expectation of all MAUDE reports is that they refer to medical devices. Indices discriminate between the cases in a generalised episode. For example, the components of Figure 14.7

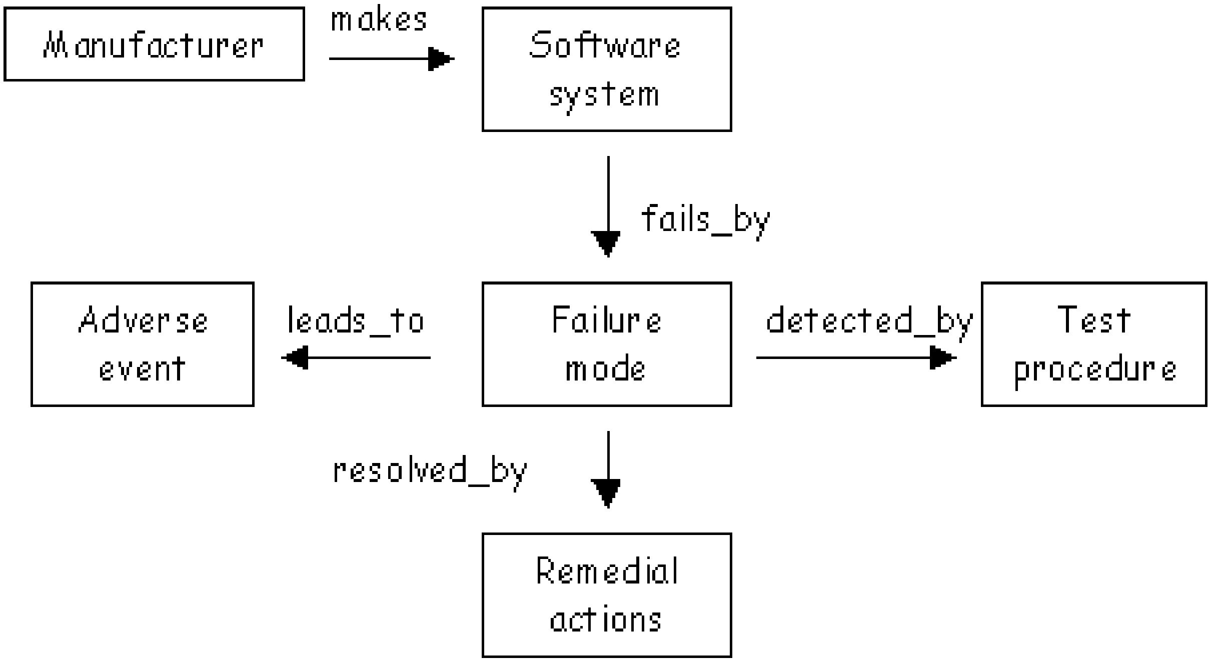

Figure 14.8: Using a Semantic Network to Model Stereotypes

might be used to index individual cases of software failure. A particular incident report could be identified by the software system involved, by the failure mode, by the detection method and so on. It is instructive to draw parallels between such system architectures and the distinction between general and particular causes that was introduced in Chapter 10.

Kolodner goes on to describe how the hierarchical structure of generalised episodes can be used to search for similar cases. The system begins at the top of the structure by examining whether or not the new incident obeys the norms associated with the episode. For instance, a retrieval task with MAUDE might begin by asking whether or not the new incident involves a medical device. If the norms are satisfied then the system examines the indices associated with that episode. These point to successively more detailed episodes. For example, a MAUDE search task might go on to consider the generalised episode associated with software failures. As before, the system examines the norms and indices associated with this form of failure until eventually an index is found that points to a matching case. The match can be computed using a 'nearest neighbour' algorithm which associates measures of similarity with each of the values that are assigned to an index. For example, if the generalised example of software failure were indexed by device manufacturer then lexical similarity might be used to identify a potential match. In this case 'Arclights technology' might return a high similarity value for 'Arclights systems' and so on. This matching process can result in the extension of the 'case memory'. If a feature of the new case matches a feature of an existing case then a new generalised episode can be created. The two cases are discriminated by creating new indices within this generalised episode. This implements a dynamic memory structure because similar parts of two case descriptions are dynamically generalised into a further episode. The significance of this should not be underestimated. Implementations of the Kolodner approach will continually update their equivalents of the semantic networks introduced in previous sections. This is done automatically as new incidents are entered into the system and hence the approach avoids many of the problems associated with 'static' data models in relational systems.

The approach advocated by Kolodner has been elaborated by a number of other researchers. For instance, the 'category and exemplar' approach distinguishes between problem descriptors and the cases that are stored in the system. Users are assumed to be looking for previous cases that describe potential solutions to the situations characterised by a problem descriptor. This approach provides three different types of indices [67]. Feature links point from problem descriptors to cases or categories. These indices are called 'remindings' because they remind users of previous solutions. Case links point from categories to associated cases. These are known as exemplar links because they indicate those cases that provide examples of the higher level category. These exemplars are ordered in terms of how well they represent this category. Finally, difference links relate similar cases that only differ in a small number of features. Unlike Kolodner's approach where there is a strict hierarchy between generalised cases the 'category and exemplar' approach uses a semantic network to link higher level categories. This supports the generation of explanations during 'knowledge-based pattern matching'. For example, Figure 14.8 supports inferences about partial reports of similar incidents. Both reports might identify the same manufacturer and the same software failure mode. If only one report named the software involved then a partial match might be made because, from Figure 14.8, the same manufacturer makes software that has previously failed in the same manner.

The previous paragraphs have describe how many features of case-based reasoning systems can be used to support search and retrieval tasks in large-scale incident collections. There are a number of further benefits [4]. For example, conversational case-based systems address the problems of poor precision and recall that frustrate the users of probabilistic information retrieval systems. In this approach, users interactively answer questions that are intended to guide them along the indices that lead to previous cases. By providing feedback about the numbers of cases that match particular responses, users can iteratively refine their search tasks in an interactive manner. For example, an initial search on the MAUDE data might prompt the user to specify who was responsible for submitting the report, whether the report addressed a hardware failure, software failure or an operator error and so on. Associated with each possible response would be an indication of the resolution provided by the question. If for example, there were only four software related incidents in the system then the user would see that by selecting this possible answer then their search would be refined down to a relatively small number of candidate cases. If, in contrast, 'operator error' indexed several thousand cases then the user could be alerted to the potential need to further refine their search task. As can be seen, this interactive approach does not directly address the underlying problems of precision and recall. The case-based reasoner may still fail to return a previous case that the user might consider to be relevant to their query. Conversely, it might return a previous incident that the user does not consider to be related to their current search task. Conversational case-based reasoning does, however, enable users to interactively control the granularity of their search task. The iterative presentation and answering of questions guides the users towards similar cases and avoids the need for users to create valid queries using a relational algebra.

The efficiency of any interaction with a case-based system can be assessed in terms of the amount of information that a user must provide in order to identify similar incidents. An inefficient system might request a mass of contextual data that does little to focus the search process. For example, a MAUDE implementation might prompt for details about high-level 'norms' within Kolodner's Dynamic Memory Model, described above. These details are likely to provide only limited benefits during any retrieval task because they will be shared by all incidents in the system. A number of algorithms exist for increasing the efficient of case-based retrieval. For instance, decision tree techniques often assign relatively high priorities to indices that partition candidate cases into a number of near equal groups. Selecting any one of the available answers will exclude a large number of cases from the other groups. If partitions are of different sizes then there is a risk that the user will continually select the index with the largest number of remaining cases and the partition will be less effective. Other algorithms have been implemented to ask questions based on their frequency of use to discriminate previous cases by other users [4].

The US Naval Research Laboratory has exploited conversational case-based reasoning techniques in the development of their Conversational Decision Aids Environment (NaCoDAE) [639] Figure 14.9 illustrates how this system supports fault-finding tasks. In this example, NaCoDAE is being used to diagnose a problem with a printer. After loading the relevant case library, the user types in a free-text description of the problem that they are faced with. The tools uses this to perform an initial search of the available cases using a form of lexical search. NaCoDAE responds with two ranked lists. The

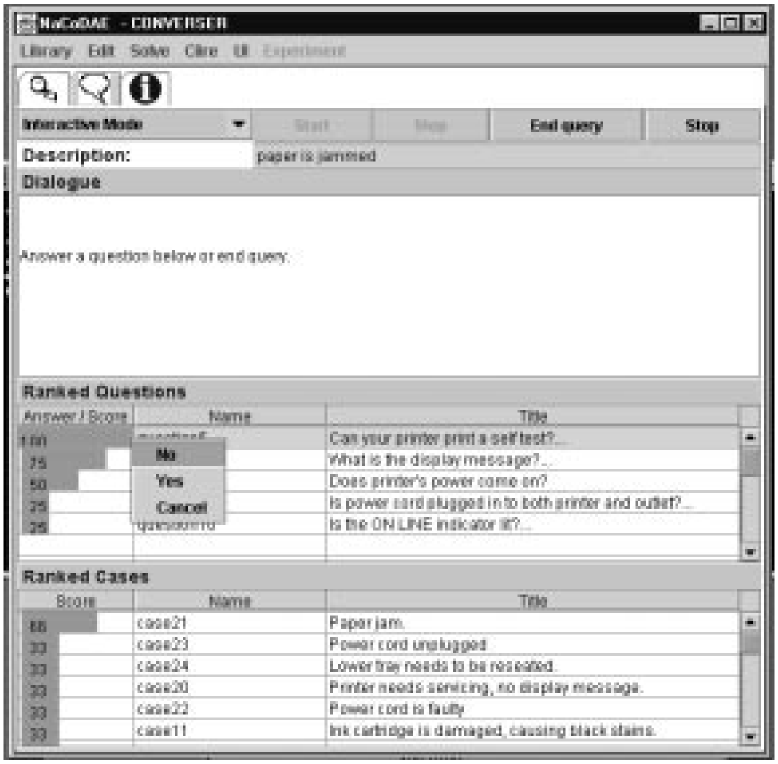

Description: paper is jammed

Dialogue

Answer a question below or end query.

Ranked Questions

Answer / Score	Name	Title
100	question5	Can your printer print a self test?...
75		What is the display message?...
50		Does printer's power come on?
25		Is power cord plugged in to both printer and outlet?...
25	question10	Is the ON LINE indicator lit?...

Ranked Cases

Score	Name	Title
66	case21	Paper jam.
33	case23	Power cord unplugged.
33	case24	Lower tray needs to be reseated.
33	case20	Printer needs servicing, no display message.
33	case22	Power cord is faulty.
33	case11	Ink cartridge is damaged, causing black stains.

Figure 14.9: US Naval Research Laboratory's Conversational Decision Aids Environment

first contains cases that are ordered using similarity measures that are based on the free-text query and the vector-based techniques that were described for probabilistic information retrieval. Each NaCoDAE case is composed of a problem description, some associated questions and, if appropriate, a description of remedial actions. The second list, therefore, presents a series of questions that are associated with the cases in the first list. The user can choose to select a possible answer to one of these questions as a means of further filtering their search. For example, they might indicate that they were only interested in cases for which their was a positive answer to the question 'was the incident detected by an end-user facility?'. The list of matching cases would then be revised to exclude those that were not detected by end-users.

In Figure 14.9, the user has typed 'paper is jammed'. The system has responded with a list of questions headed by 'Can your printer print a self-test'. As mentioned, this question guides the user in their retrieval task. If they did not understand the question then they can double click on the question to reveal a further explanation:

> "To perform a self-test, make sure that the printer is off-line and while holding down the ALT key, click the TEST button'.

If the user cannot follow these instructions they can continue the search by answering another question from the list. In Figure 14.9, the user has indicated that the self-test procedure failed. The cases displayed below can then be revised in the light of this additional information. This co-operative exchange of questions and answers will also help improve recall because the user can continually review the list of 'relevant' cases being retrieved at each stage of the process. If the user selects the 'Paper jam' (Case 21) then they will receive further information on corrective actions. The information encoding used by NaCoDAE can be illustrated by this example:

```
BEGIN QUESTION QUESTION5
   TITLE 'Can your printer print a self test?...'
   TEXT ‘To perform self test, make sure printer is OFF-LINE, and while
      holding ALT key, click the TEST key.’
   ANSWERS
      TYPE : YES_OR_NO
   WEIGHT
      MATCH : 10
      MISMATCH : 2
   AUTHOR david_aha
   CREATION DATE 7/30/91 TIME 15:18:33
   LAST_UPDATE DATE 7/30/91 TIME 15:18:34
END QUESTION
```

As can be seen, the initial self-test question includes information about how to reach a potential answer. It also states that the type of the answer must be a YES_OR_NO. The run-time environment provided by the case based reasoning tool interprets this information and presents the user with a drop down menu which constrains them to a ‘yes’ or ‘no’ answer. The weighting information can be used in a variety of ways. The simplest approach is to increment the weighting of any cases matching the users' selected answer and a penalty for cases that do not match the selected response. The encoding also includes information that supports the maintenance of a case base by denoting the identity of the person who entered the question into the case-base and the date of last modification. Individual cases can be encoded in a similar fashion. As can be seen, the developer explicitly states the responses to particular questions that will increase the weighting associated with a particular case. In this instance, if the user selected ‘no’ in response to question 5 ‘Can your printer print a self test?’ then the match would be incremented in the manner described above. Conversely, if a self test was completed then the weighting would be decremented:

```
BEGIN CASE CASE21
   TITLE ‘Paper jam.’
   QUESTIONS
      Question5 :  ‘No’ (MATCH_WEIGHT : + MISMATCH_WEIGHT : -)
      Question25 :  ‘13 Paper Jam’ (MATCH_WEIGHT : + MISMATCH_WEIGHT : -)
   ACTIONS Action23
   CREATION DATE 8/15/91 TIME 10:56:51
   LAST_UPDATE DATE 8/29/91 TIME 18:42:1
   LAST_USED DATE 8/15/91 TIME 10:56:51
   NUMBER_OF_CALLS 0
END CASE
```

NaCoDAE's encoding of individual cases identifies potential solutions. The paper jam case number 21, illustrated above, is associated with remedial action number 23. This is represented by the following formalisation. As can be seen, the proposed intervention is identified by a short title ‘Clear paper path and reseat paper cassette...’ as well as a more sustained series of instructions. These end with a final recommendation that if the problem persists, users should contact a service engineer:

```
BEGIN ACTION ACTION23
   TITLE ‘Clear paper path and reseat paper cassette...’
   TEXT ‘Jamming can be caused by crooked cassette, wrong paper type,
      wrong side of paper up, or sometimes by a dirty print bar or worn
      tractor wheels.  If the problem persists, contact a service
      representative.’
   AUTHOR david_aha
   CREATION DATE 8/15/91 TIME 10:56:40
   LAST_UPDATE DATE 8/15/91 TIME 10:56:42
END ACTION
```

Previous sections have described how NaCoDAE represents each case in terms of a free-text description, a set of appropriate actions and the answers to questions that help to classify the case. The MAUDE data-set readily provide descriptions for each incident in the form of the free-text reports that were associated with each record. MAUDE does not provide access to detailed information about the response to individual incidents. We cannot, therefore, directly encode MAUDE records within NaCoDAE. Our initial studies overcame this problem by referring the user to a range of documents provided by the FDA about appropriate responses to general types of device failures, including recall and emergency response guidelines [257], reporting delegation procedures [252] and risk management documents [276].

It is harder to identify appropriate questions that might be used both to partition the data set and to guide the users' search. The most straightforward approach is to derive questions directly from the existing relational data model. For instance, users might begin a search by answering the question 'what was the outcome of the incident?'. They could then select an appropriate response from the alternative answers 'death, injury, malfunction, other'. This might result in the retrieval of a number of cases that contained either positive or negative answers to the question 'was the anomaly reported by a manufacturer?'. This question would then be presented to the user as a way of further refining their search using information that was common to the cases from their initial query. All of this data is readily identifiable from the existing MAUDE database. However, it is important to ask whether this encoding would offer any benefits over the traditional database approach? The first benefit is that NaCoDAE does not associate an answer for every question with each case in the system. This is appropriate because, as we have seen in Chapter 5, there can be considerable uncertainty about the causes and consequences of some incidents. For instance, the person submitting the form may not know how it was resolved. NaCoDAE actively exploits the absence of information because it helps to distinguish between different cases. If a user decided not to answer a question then their search will retain cases with these 'unknown' values. However, if they select a definite answer then these cases will be excluded along with cases that are associated with the alternative answers to that question. In contrast, most relational implementations specifically prohibit 'absent values' from the fields of a record. Many relational systems, therefore, resort to using 'other' as a potential value that can be recorded. This does not resolve the problem, however. There is an important distinction between other' which implies that a definite response is known but is not supported by the system and 'unknown'. Some relational systems support the distinction by including both 'other' and 'unknown'. Unfortunately, this creates frequent problems during the training and appraisal of coders who must be reminded of the difference between these two potential values.

The previous paragraphs have described how our initial application of NaCoDAE was restricted to the information that was included in the original MAUDE reports. The relational data model that supports the existing database only provides limited causal information. Most of this detail is embedded within the natural language accounts that are associated with each incident report. Previous sections have argued that lexical retrieval techniques can be used to identify common features in the language that is used in these accounts. Unfortunately, there is no reliable automatic means of extracting causal information from these natural language accounts. We, therefore, decided to recode our MAUDE sample data to demonstrate that case based reasoning tools, such as NaCoDAE, can be used to support the direct search for common causal factors. This builds on previous work in the application of case-based reasoning to 'small-scale' incidents by Koornneef [456]. The first stage of this new work was to perform a causal analysis of the incident reports. This followed the Eindhoven classification technique described in Chapter 11 [840, 844]. The causal analysis associated each incident with a number of the leaf nodes shown in Figure 11.10. Fr instance, the following natural language description provides an informal account of the potential causes of an adverse event involving an insulin infusion pump:

> "Patient treated at hospital for hyperglycemia. Pump not returned for evaluation... Manufacturer could not evaluate the pump, as the patient did not return it. User error likely caused event. Continuous insulin infusion therapy requires that the patient continually assess the impact of such factors as their caloric intake, activity levels and other medical conditions and/or treatments on their blood glucose level. The ther-

> apy also requires periodic self-testing of actual blood glucose levels. Failure to monitor and/or adjust the insulin amount appropriately will result in erratic blood glucose levels. Extreme excursions from normal blood glucose levels can result in conditions such as hypoglycemia or hyperglycemia. Patients experiencing these conditions may require hospitalization and medical intervention to preclude serious medical conditions including death."

Previous chapters have identified a number of criticisms that might be made both about the style and content of this account. Strong assumptions are made about the patient's role in the incident. For instance, we are told relatively little about the information and guidance that the physician offered to support the use of the device. Chapter 11 has presented techniques that can be used to address these concerns. For now, however, it is sufficient to observe that the causal factors related to this incident might be categorised using the HRM (Human Behaviour: Monitoring) and PRF (Patient Related Factor) nodes from Figure 11.10. The incident was caused by a failure on behalf of the patient and clinician to monitor their use of the device. The incident was also caused by specific patient related factors, including their underlying medical condition that led to the hyperglycemia. We then encoded this analysis as positive responses to the questions 'was there a failure in human monitoring?' and 'was the incident exacerbated by patient related factors?'. Conversely, there was no evidence of a device related failure (TD, TC or TM). This was encoded as negative responses to the questions 'was there a problem with the device design?', 'was there a problem with the device construction?' and 'was there a problem with the device materials?'. As mentioned, we did not have to indicate whether or not each element of the Eindhoven classification was a causal factor for every incident. Answers were only provided when there was definite evidence for or against certain causal factors. Instead of questions about the facts known for each incident, such as the name of the device or the manufacturer, these changes support the classification of cases or incidents by the results of the causal analysis. Not only does the NaCoDAE application support direct queries of the form 'who reported the incident?' but it also supports searches that look for complex combinations of causes such as 'what incidents were not reported by manufacturers but were caused by a lack of monitoring on behalf of the device user or clinician?'. Such queries cannot easily be satisfied using conventional databases and information retrieval engines.

The previous paragraphs have described initial attempts to apply case-based reasoning as a partial solution to the problems identified for relational databases and lexical information retrieval systems. As we have seen, however, many case-based systems draw upon ideas that were originally developed to support these more common applications. For instance, NaCoDAE relies upon a lexical analysis to perform the initial identification of candidate cases and questions. Similarly, the semantic networks of many case-based systems can be thought of as dynamic versions of the data models that underly relational systems. There are also strong differences. In particular, most case based systems do not require an initial domain model. The classification emerges over time as new cases are added to the system. This is a significant benefit given the widespread disagreement that exists over appropriate incident classification schemes [417].

A number of further issues must be addressed before case based reasoning techniques can be widely applied to support the storage and retrieval of incident reports. For instance, it is unclear how to provide the system with feedback when users disagrees with the matches that are proposed for particular incidents. This is complicated because such matches often depend on indices that have been automatically inferred by the system. It is, therefore, important to provide the user with information about the reasons why the system identified a target incidents as being similar to the one under consideration. Some systems address this issue by simply showing the user a trace of all of the factors that match between the situation that they are describing and the one that has been retrieved. Under such circumstances, the user can then either revise their interaction with the system or alter the labels associated with the case that was erroneously retrieved. The user might provide additional indices to distinguish the new incident from the case that was incorrectly matched. This can create problems if arbitrary users are permitted to alter the indices that are generated by the system. Different users are likely to disagree about the appropriateness of a particular match. It is for this reason that NaCoDAE records authoring information with the insertion of new cases and questions into a case library.

This chapter has focussed on the use of case based reasoning systems to identify patterns during the retrieval of incident reports. There is a danger that this focus will obscure the main motivation behind the development of this technology. Case-based reasoning systems were originally intended to help users solve problems and make decisions. Most previous applications of this technology, therefore, also include some assessment of how effective a proposed intervention was in response to previous cases. We have not been able to introduce this information into our initial studies using the NaCoDAE system because MAUDE does not assess the effectiveness of interventions following individual device failures. If this data were to be made available then regulators and analysts could use the case-based retrieval facilities of NaCoDAE to ensure that they respond to situations in a consistent manner. Users could also determine the circumstances in which a particular intervention had previously been effective. Without such assistance, there is a danger that the system would consistently advocate the wrong intervention. The next chapter, therefore, focuses on techniques that can be used to monitor the effectiveness of incident reporting systems and the recommendations that they produce.

14.5 Summary

Previous chapters have described the elicitation and investigation of adverse incidents and near-miss events. We have also considered a range of different techniques for presenting the findings of these investigations. In contrast, this chapter has looked at the issues that arise when regulators and safety managers must disseminate information about these safety-related occurrences. It is important not to underestimate the scale of this task. For example, the UK MDA provides information on approximately 7,000 incidents each year [539]. The US FDA's MedWatch program generates well over 300 incident-related publications each year [269]. The tasks associated with disseminating this information are exacerbated by the tight deadlines that must be met if safety managers are to be provided with the information that is necessary to respond to adverse vents in a timely manner. The MDA has a commitment to issue Hazard Notices within 20 days of notification, safety notices should be issued within 90 days. There are also financial pressures. The MDA are expected to meet these targets while at the same time achieving 2% efficiency savings per annum.

The pressures of time and of economy have led many reporting agencies to carefully consider who should receive the information that they disseminate. Some systems operated closed distribution policies where reports are only passed to a few named individuals within an organisation. Horizontal systems distribute to safety managers within other companies in the same industry. Vertical distribution schemes disseminate reports widely within the same company. Parallel reporting systems distribute reports to companies that operate similar processes in a range of different industries. Open distribution policies place few restrictions on the recipients of incident reports. Although we have identified these general approaches, many organisations operate hybrid techniques. For example, the MDA distribute Safety Notices through the Chief Executives of Health Authoritities, NHS Trusts and Primary care Trusts as well as the directors of Social Services in England. This represents a parallel approach to dissemination because each of these individuals may be responsible for similar healthcare systems that operate in very different contexts. However, each of these individuals is then responsible for further disseminating the Safety Notices widely to 'all who need to know or be aware of it' [535]. Hence this second stage dissemination opens up access to a far wider audience.

The carefully designed distribution policies that have been devised by many reporting agencies are often undermined by alternative communication channels. For example, informal anecdotes and 'war stories' provide both a powerful means of self-help and a dangerous source of rumour depending on the information that is conveyed and the context in which they occur. These informal channels are becoming increasingly important as technological innovation is increasing individual access to wider distribution media. In particular, the development of Internet chat rooms and of 'special interest' web pages has led to the dissemination of many 'alternate' accounts for incidents and accidents. The press and broadcast media provide further means of disseminating information about adverse events. They may be used to publicise the findings of an official investigation. They can also disseminate the results of journalistic investigations which are, typically, triggered by members of staff who feel that safety-related information must be disseminated to a wider audience.

Incident reporting agencies can recruit a range of technologies to implement the distribution policies, described above. These range from conventional paper-based publications through to increasingly complex, computer-based storage and retrieval systems. Paper based resources have numerous benefits. They are accessible to a wide audience and impose few additional technological requirements either on the publishers or the recipients. Unfortunately, they can be costly to produce and are difficult to disseminate in a timely fashion. There are also limitations in the types of information that can be captured in books, pamphlets and journals. Some information can be better conveyed using more dynamic media such as video images of incident locations and computer reconstructions of likely events. Finally, it can be difficult to ensure that all readers receive a copy of periodic updates to paper-based incident reports. There is a danger that some safety-managers may retain printed documents that contain obsolete recommendations.

Many of the distribution problems that are associated with paper-based documents can be addressed through the use of fax and telephone based mechanisms. For instance, pre-recorded messages can be accessed by telephone so that the potential recipients of an incident report can determine whether or not to request a printed copy using more conventional means. Alternatively, fax machines can automatically send updates to many thousands of telephone subscribers. This can be done over-night or during periods when the necessary equipment is likely to be idle. Unfortunately, the low resolution of most fax devices and the relatively unreliable infrastructure can create problems if these approaches are used as the primary means of dissemination. Increasingly, reporting agencies view this form of technology as an interim measure while the intended recipients of their documents acquire the necessary support to access computer-based resources.

A range of issues must be considered by any organisation that is considering using computer-based systems as a means of disseminating safety-related information. They must consider whether machines will be isolated from the security concerns that are associated with many local and wide area networks. They must consider whether the information that is held on a machine is to be disseminated by transient media, such as email, or by more durable forms of secondary storage, including CD-ROMS. They must consider the way in which individual reports will be formatted. For example, variants of the HyperText Markup Language (HTML) and Adobe's proprietary Portable Display Format (PDF) are both emerging as standards for the transmission of incident reports over the web. Each of these approaches offers radically different support for the dissemination of incident reports. PDF provides better support for the local generation of printed documents. HTML is more easily indexed and searched by a wider range of automated systems. Investigation authorities must also approve the access control mechanisms that are intended to secure their information resources. They must consider who has the right to read an incident report. They must also consider whether those initial readers have the right to disseminate the report more widely once they have received it. These access control mechanisms must be implemented using techniques such as public and private key cryptography. These techniques can be used to establish that the information has been sent by an official source and that the recipient has the correct permissions to access any data. Digital watermarks can also be used to ensure that incident information has not been altered by a third party. Finally, reporting agencies must also consider accessibility issues. Many schemes operate within regulatory and legal frameworks which help to ensure that the use of particular technologies does not prevent potential recipients from reading an incident report. Most commonly this is interpreted as a requirement to provide information in a format that can be accessed by individuals with a visual disability, for instance using a screen reader. Some legal and regulatory requirements have wider implications, including the need to perform usability evaluations to establish that computer-based resources can be operated by a wide cross-section of potential users [686].

These issues are generic because they affect the application of computer-based technologies to support the dissemination of incident-related information. There are, however, a number of more specific concerns that stem from the use of particular computational techniques in this domain. For instance, most existing systems rely upon relational databases. These applications structure the storage and retrieval of information using a static data model that must be carefully designed before the system is implemented. These models can be refined to improve efficiency both in terms of the storage space that may be required and in terms of the access speed for individual reports in large scale systems. Relational data models also help to ensure that data is not omitted or needlessly repeated. There are further benefits. Relational data models associate particular fields of information, or attributes, with key entities in the application domain. For instance, the FDA's MAUDE system is structured around manufacturer, device and patient 'entities'. The attributes associated with these key entities help to define the minimum information that must be recorded about each incident. This increases consistency between individual incident reports. Unfortunately, many of the benefits of relational databases can also be interpreteted as potential weaknesses. For instance, these is often considerable confusion about the values that must be entered into the individual fields of a relational system. The MAUDE system supports the distinction between a 'generic name' and a 'brand name' that can be confusing without further explanation. Such problems can also frustrate information retrieval using relational systems. Queries must, typically, either be pre-formatted or composed using a variant of the relational algebra. If queries are pre-formatted then it can be difficult for designers to anticipate all of the questions that users might need to pose of the incident data that is collected. If 'raw' queries are to be constructed from the relational algebra then users must not only be very familiar with the underlying data model but they must also have some understanding of the particular operators supported by their database management system.

A final set of limitations stem from the static nature of many relational schemas. As mentioned, most database applications structure the storage and retrieval of infromation around a number of tables or relations that are 'optimised' to improve the efficiency of a resulting application. These tables must be 'pre-programmed' into the system. In consequence, it can be difficult to develop appropriate models if safety managers or regulators are unsure about the precise nature of the incidents that will be reported or the information that they wish to capture. This might seem like a trivial requirement; the operators of a reporting system should have a clear idea of the information that they wish to elicit before starting a scheme. Unfortunately, things are rarely this simple. Even if it is possible to identify information requirements before a system is established, those requirements are highly likely to change over time. For instance, changes in production techniques may lead to new questions being asked about the circumstances in which an incident occurred. This would force programmers to refine the attributes in a relational data model. Similarly, if a relational model were to include causal information then changes might have to be made whenever new causal factors were identified. This would raise particular problems if previous incidents in the database were not reclassified using the new causal model. For instance, if a system added 'high workload' as a new cause in January 2002 then all the system would only recall incidents after this date even though there may have been 'high workload' incidents received before this date. These earlier failures would not have been recorded in this way because the relational model operating before 2002 did not support this distinction. For systems, such as MAUDE, that contain several hundred thousand records the maintenance issues associated with relational systems can impose a considerable overhead upon systems administrators.

Many incident reporting systems avoid the limitations of static relational models by restricting the information that is encoded within the fields of the database. These fields only record information about entities that will not change over the lifetime of the system. Every medical device will have a manufacturer, every incident report will be submitted by a contributor and so on. The remainder of the contextual and causal information that will change over time is recorded in a free-text description of the adverse event. This approach is adopted by the MAUDE system. Free text information retrieval offers numerous benefits for the maintenance of large-scale incident reporting systems. Many of these techniques do not assume that users have any knowledge of the underlying implementation techniques. Nor do they require the use of complex relational algebras. In contrast, users are encouraged to form natural language queries. Typically, key terms are extracted from these queries. There terms are then compared against the indices that point to individual narratives. Stemming techniques and thesauri can be used to ensure that lexical retrieval systems detect matches even though users do not enter exactly the same literal terms that are indexed by the system. Hence, 'fail', 'failed', 'failure' and 'fallible' can all be recognised as referring to similar concepts. If there is a sufficient match between the terms in the query and the index terms in the document then the system will propose a potential match to the user. This lexical retrieval clearly depends upon there being a minimal distance between the language used in the query and the index terms. If users' continually uses words that the system does not recognise then it will be difficult to identify appropriate documents.

Expert advice can be used to guide the selection of appropriate indices. This advice can be validated against records of queries performed with previous versions of the system. Index terms can be identified by a lexical analysis of word frequencies. This has the benefit that index terms can be revised to reflect changes both in the language that is used to describe incidents and, ideally, in the nature of the incidents themselves. Any changes in word frequency will be accounted for each time the indexing program is run. One side effect of this is that free text retrieval systems avoid the limitations of more static relational schemes. A potential problem with this approach is that requests will not always return the same results because the indices depend on the changing frequency of terms used in the collection.

Further problems arise because it can be difficult to ensure both high precision and high recall over a broad range of user queries. Precision refers to the proportion of relevant to irrelevant documents that are returned in response to a query. Recall refers to the proportion of relevant documents that are returned to all relevant documents held in a collection. Both of these concepts are well illustrated by the current generation of web technology. Many search request now provide hundreds of potential 'hits'. Many of these will not be relevant to the users query. The manual process of sorting through these many irrelevant matches stems for poor precision. Similarly, the same request may not return all of the potentially relevant information. There may be sites that could have provided exactly what the user required but which were not recognised as being relevant. This illustrates poor recall. In the context of incident reporting, each of these concepts has considerable significance. A failure to retrieve a similar incident in the past may mean that safety managers fail to detect an emerging pattern of failure. Poor recall can, therefore, lead usrs of a system to underestimate the potential risks of any recurrence. Similarly, if a request returns many dozens of incidents that the user does not consider to be relevant then they may be dissuaded from performing the necessary manual filtering that might have identified previous similar incidents. Such poor precision can impose considerable burdens upon the finite resources of many incident investigation agencies.

A range of solutions have been proposed to avoid the limitations of relational databased and lexical information retrieval systems. As mentioned, the FDA's MAUDE system implements a hybrid interface that provides access both to a relational database for directed search and a free-text retrieval system for broader queries. Case-based reasoning tools have recently been identified as a further alternative [456]. These systems, typically, avoid any predetermined data model. Instead, they will automatically reconfigure indices as new cases are entered into the system. There are several ways of achieving this. For instance, the Dynamic Memory Model distinguishes between 'generalised episodes' that collect together similar cases and indices that are used to distinguish between each of the particular cases that represent instances of a 'generalised episode'. Alternatively, 'category and exemplar' approaches introduce several different types of indices some of which indicate the degree of closeness between a higher level category and the incident reports that are exemplars of that category. These approaches offer significant benefits to incident reporting systems. Not only are they based upon dynamic classification techniques, most of these tools have been deliberately designed as information support systems. Individual cases are usually formulated as descriptions of problems. These are then associated with remedial actions. Hence it is possible to find out what other circumstances might prompt similar interventions. It is also possible to determine whether the same incidents are provoking the same reaction. These standard features of many case-based reasoning tools must be explicitly designed into mass-market relational databases.

The chapter has closed by describing initial attempts to apply the US Naval Research Laboratory's Conversational Decision Aids Environment (NaCoDAE) to store and retrieve reports of device failures from the MAUDE collection. NaCoDAE is a conversational case-based reasoning tool. Users provide an initial free-text query. This is used to identify an initial set of matching cases. This initial match is then analysed to identify a series of questions that might best be used to discriminate between these individual incidents. The user is then prompted to answer a list of these questions. For example, the user could choose to select the answer 'yes' to the question 'was the incident reported by an end user facility'. Each time they select a response, the system will automatically revise the set of matching cases and the list of questions. The efficiency of the entire system can be judged in terms of the number of questions that must be answered before the user is satisfied that they have identified a potential match. This ‘conversational’ approach helps the user with the problems of query formation. They are continually prompted to answer questions that are intended to guide their search task. This approach also helps the user to control the number and nature of potential matches. It is a trivial task to ‘cancel’ a response to a question if it narrows the number of potential matches too rapidly.

As mentioned, an important benefit of case-based tools is that they explicitly support the association of incident descriptions and recommended remedial actions. They do not, however, guarantee that those remedial actions will either be effective or appropriate. The following chapter, therefore, described ways of monitoring the effectiveness of the interventions that are identified in response to adverse events and near-miss incidents.

Chapter 15

Monitoring

The New Zealand Transport Accident Investigation Commission argues that “the sole purpose of each (incident) report is to avoid similar occurrences” [632]. This chapter looks beyond this high-level goal to identify the problems that arise when attempting to monitor the success of any incident reporting system. For instance, it can be difficult to prevent similar occurrences when new technology and working practices introduce new froms of previous failure. It can also be difficult to ensure that any reporting system gathers sufficient information about other failures to be sure that similar failures are not going unreported. Given such uncertainty, it is particularly important that any monitoring activity justifies the investment that regulators and operators must make to sustain reporting systems. Previous chapters have described the many different activities that must be managed during the investigation, analysis and dissemination of incident reports. Domain experts must initiate follow-up interviews, site visits and data acquisition. They must filter relevant information from the mass of contextual details that are elicited in the aftermath of an adverse event. They must also ensure that the products of any root cause analysis are well-documented so that others can reconstruct the arguments that support particular recommendations. We have seen that computational tools can assist in the elicitation, classification, dissemination and retrieval of incident reports. Computer-based forms can be developed to collect initial information about an adverse event. Automated interviewing systems can prompt domain experts to consider certain causal hypotheses in the aftermath of an incident. The previous chapter has described the use of lexical retrieval tools, relational databases and case-based reasoning systems to identify patterns of failure. It is important to emphasise, however, that these tools have not been widely applied to support incident reporting systems. The costs of manually performing these various activities can, therefore, act as a significant disincentive to the creation and maintenance of many reporting schemes. It is, therefore, important that safety managers can demonstrate the ‘cost-effectiveness’ of any proposed system.

Who Monitors What and Why

A host of problems complicate attempts to monitor the ‘cost-effectiveness’ of incident reporting systems. For instance, it can be difficult to establish that any safety improvements can be attributed to a reporting system rather than to other wider changes within a company or an industry. A more detailed examination of the barriers to incident monitoring is postponed until later sections. It is first important to identify those individuals and organisations that are concerned to validate the effectiveness of a reporting system.

Safety Managers. There are many different reasons for monitoring incident data. For instance, safety managers can use information about adverse events to justify the remedial actions that are intended to prevent future failures. This also, indirectly, helps to justify the existence of the reporting system. Safety managers can also monitoring incident reports to identify the need for further contributions about certain safety concerns [169]. They can use incident information to track progress towards higher-level safety targets. They can also use monitoring data to inform employees and the public about particular safety issues and so on. For example, the Washington Metro Area Transit Authority's Safety Officer issued a press release to publicise customer injuries on escalators [855]. These accounted for 43% of all passenger injuries in the first quarter of 2002, a 5% reduction from 2001. The Safety Officer argues that this represents significant progress. He also uses this data to justify a new safety awareness campaign involving community outreach, new car cards, station posters and brochures. This example illustrates the way in which incident monitoring helps Safety Officers to recognise the scale of particular problems. They can then use the available statistics to inform the public about potential hazards. The same press release also illustrates the way in which incident monitoring can be used to provide feedback to employees about previous safety initiatives. The Safety manager describes how there were 13 reported fires and smoke incidents requiring a fire service response in the first quarter of 2001. These resulted in an average delay of under 30 minutes per incident. In the first quarter of 2002, there were only 10 such incidents with an average delay of under 20 minutes. The Safety Manager concluded that 'the reduction in the number of incidents is a result of improved maintenance measures, and interagency coordination and communications with the local fire departments' [855].

The information provided by the Washington Metro Area Transit Authority illustrates the use of incident monitoring to track relatively long term trends. Safety managers can also use this data to identify sudden increases in particular types of failure. For example, Southern California's Metrolink Rail System has used monitoring information to provide a rapid response to changes in the types of incident that are being reported [548]. In February 2000, it was realised that there had been 5 different incidents involving trains and trucks in the Southland area in a 90 day period. This formed a sharp contrast with the previous 24 months in which there had only been a single comparable incident. The five more recent incidents were 'near misses' in the sense that they resulted in relatively minor injuries. The careful monitoring of these incident statistics helped to trigger a more detailed causal analysis. This identified that a booming regional economy had resulted in an increase in freight carrier traffic. Many of the truck drivers who were brought in to satisfy this demand were unfamiliar with the Metrolink train operations. The company responded to this rapid increase in truck-train incidents by contacting local haulage firms and by a series of awareness raising initiatives including a 'Trucker on the Train' day as part of a Metrolink Rail Safety Week. This enabled truck drivers to ride with Metrolink engineers so that they can learn to avoid potential collisions.

The two previous examples have focussed on the use of incident monitoring to chart progress towards recognised safety objectives and to trigger rapid intervention when new hazards arise. This illustrates how safety managers can use data to provide others with information about adverse events and near miss incidents. It is also important to monitor the reporting system itself to ensure that submissions are handled in a timely fashion. Safety managers and the operators of reporting systems must also track any causal analysis to identify potential bias [854]. This must include some consideration of intra-rater reliability; will the same coder code similar incidents in the same way over time. Similarly, the findings of particular analysts may be compared for the same incidents to ensure inter-rater agreement. Davies, Wright, Courtney and Reid describe the results of performing this type of monitoring activity for the CIRAS voluntary reporting system that operates across Scottish railways [197]. This system is based around a classification system, similar to those described in Chapter 11, which provides 54 different causal categories. Two CIRAS personnel independently analysed a total of 439 incidents with 84.6% agreement over the causal classification. This involved the assignment of 1,955 codes for human factors issues alone. Such results represent a remarkable level of consistency. Chapter 12 has reviewed the many problems that can jeopardise agreement between independent analysts. For example, Lekberg has shown that individuals from different educational and operational backgrounds will code the same incident in different ways [484]. Monitoring techniques, such as those introduced in the following pages, provide Safety Managers with a means of assessing whether or not such factors are introducing significant biases into the analysis of and response to adverse events.

Company Monitoring. The previous section described some of the reasons why Safety Managers might choose to monitor the performance of an incident reporting system. At a corporate or organisational level, there can be more pressing requirements to track the data that can be obtained about incidents and accidents. For example, Canada's Railway Safety Act incorporates an annex that describes various requirements that must be satisfied by Railway Safety Management Systems. These include the provision that all railway companies must record safety-related information for the purpose of 'assessing its safety performance' [779]. This information should include 'accident and incident investigation reports and a description of the corrective actions taken for accidents and incidents'. Companies must also monitor accident and incident rates that should be calculated in terms of (i) employee deaths, disabling injuries and minor injuries, per 200,000 hours worked by the employees of the railway company, and (ii) train and grade crossing accidents per million train miles. The Railway Safety Management Systems annex also states that railway companies can be required to collect, maintain and submit specified performance or safety data for "the purpose of monitoring the effectiveness of its safety management system and its safety performance". These are important provisions because they specify the way in which companies must normalise their incident data to account for differences in the operating characteristics of individual companies. Clearly, raw incident frequencies for national carriers and for local railways cannot provide an adequate means of comparison. The number of journeys as well as the distance and time of travel combine to make the risk exposure radically different in each of these cases. It is possible to have a profound effect on the nature of safety statistics depending in which of these normalising factors are used. In contrast to the Canadian provisions mentioned above, the FRA calculates the total accidents and incidents rate by multiplying the number of accident and incident reports by 1,000,000 and then dividing the result by the sum of train miles and hours. This reflects a different approach in which companies do not directly perform the normalisation themselves. This is done by the regulator in assessing the performance of each operator. Such an approach raises a number of dilemmas for operating companies that must monitor accident and incident rates. In particular, they must still report normalisation statistics for periods in which they may have few or no safety related occurrences. This is necessary if regulators are to assess the performance of an industry in terms of total incidents per miles travelled, passengers carried etc. The government of South Australia has recently eased the burdens associated with the reporting of normalising factors through the development of a web page [782]. This asks operators to report how many kilometres of track they own and manage within Southern Australia. They should also report the distance, in kilometers, that their passenger or freight trains travelled within the state. This distance must be distinguished from the kilometers travelled by contract services. In addition, operators must report the number of passenger journeys in urban areas within Southern Australia given as 'a point to point journey irrespective of the number of vehicles or mode used for the trip'. Journeys in non-urban areas consist of 'a point to point journey but each change of vehicle along the route is a separate journey'. Companies must also report their total number of employees engaged in railway work in South Australia. This includes contractors and volunteers who work 'at the direction of the reporting railway' but not 'employees, contractors or volunteers of other accredited railway owners or operators who provide services to your organisation'.

Legal and regulatory provisions are not the only reasons why companies may monitor an incident reporting system and the data that it provides. There may also be strong commercial motivations. For instance, incident data is often cited when two or more companies are in competition for a particular market. Similarly, incident information will often be published if the operational activities of a company are called into question by the public, press, politicians or other pressure groups. For example, the San Jacinto Rail company is in the process of applying to transports hazardous and non-hazardous materials in the Houston area. Approximately 85% of the materials to be carried by the proposed service will be both solid and non-hazardous including polyethylene and polypropylene plastic resins. The proposals also provides for the transportation of more hazardous commodities, including isobutylene, propyleneglycol and ethylene glycol. In order to reassure potential opponents to this proposal, the rail company cited incident statistics gathered by a range of trade organisations:

> "According to research by the Association of American Railroads (AAR), 99.996% of hazardous materials moved by rail arrive at destination without incident. Over the past 20 years U.S. railroads have invested in technology and infrastructure to improve safety, reducing accidents per million train miles 66% since 1980 and 18% since 1990. Although trucks and railroads carry almost the same amount of hazardous materials, the trucking

> industry has nearly 14 times more hazardous material incidents." [723]

This quotation illustrates the way in which individual companies can draw upon the incident and accident data that is collected by trade associations. In this case, the San Jacinto proposal exploits the results of the AAR monitoring to compare the safety record of rail transportation with that of the haulage industry. The proponents of this scheme also cite incident data from similar companies. For example, the Burlington Northern and Santa Fe Railway experienced 17 releases of hazardous materials from a total of 878,428 shipments in 2000. The proposers calculate that this represents an 'accident release ratio' of 0.0194 accident releases per 1,000 shipments. This represents a decrease from 0.0201 accident releases in 1999.

The Nuclear Energy Institute provides a further example of how incident monitoring information can be used to support particular commercial operations [381]. A recent report stated that the US nuclear energy industry has completed more than 3,000 shipments of used nuclear fuel covering 1.7 million miles over the last 35 years without any injuries, fatalities or environmental damage from the radioactivity of the cargo. This is an interesting argument because it reverses the usual claim that reporting systems provide important information about previous failures. In contrast, the Nuclear Energy Institute stress the absence of incidents in order to reiterate the industry's safety record. This style of analysis can seem complacent. The Institute is, however, careful to stress the more active safety measures that protect the public 'accidents can happen and so scientists and engineers designed used nuclear fuel shipping containers to be among the safest on the road, and to protect the public against even the most unlikely accidents' [381].

Regulatory and Investigatory Oversight. Governments are concerned to both promote and ensure the safety of national industries. They, therefore, direct industry regulators to gather a range of statistics to monitor the performance of those industries. Some of these indicators are relatively easy to obtain. For instance, Table 15.1 presents Australian Transport Safety Bureau (ATSB) and Australian Bureau of Statistics data on fatalities in each of the major transport modes over the last decade [51]. This shows a reduction in the total number of fatalities across all modes except air transport. However, the periods in Table 15.1 reveal the lack of national, annual fatality statistics for particular industries. The lack of reliable statistics is worse for less serious incidents [51]. The Australian government has particular problems in gathering information about serious road injuries. This partly arises from the inconsistent definitions about what is and what is not reportable at a state level. Similar problems in the rail industry led to calls for a national coordinating body to receive and review incident and accident statistics [55]. The need for such a body is illustrated by the diverse legislation that covers the Australian national rail system. For example, Victoria follows a Transport (Rail Safety) Act of 1996 and enforces Transport (Rail Safety) Regulations proclaimed in 1998. Western Australia follows a Rail Safety Act of 1998 and Rail Safety Regulations of 1999 New South Wales introduced a Rail Safety Act in 1993. Section 44 of this act was affected by the Administrative Decisions Legislation Amendment Act of 1997. It also enforces Rail Safety (Offences) Regulations of 1997.

	Road	Rail	Water	Air
1991-92	2084	54	69	46
1992-93	1874	55	69	63
1993-94	1994	37	71	63
1994-95	1984	55	58	56
1995-96	1986	37	51	68
1996-97	1873	38	55	34
1997-98	1768	42	47	51
1998-99	1774	..	..	44
1999-00	1783	..	..	45
2000-01	1775	..	..	57

Table 15.1: Australian Transport Fatalities by Mode (1991-2001)

Collating statistics from different local and regional reporting systems is one of several problems that complicate the regulatory monitoring of particular industries. Chapter 2 has described the problems that arise when regulators become involved in both the promotion and monitoring of a safety-critical applications. There may be a temptation not to publicise adverse statistics that might affect the future success of commercial organisations. This explains why many countries deliberately separate the promotion and safety regulation of their industries. Even so, there is a temptation for regulators to focus on those statistics that illustrate the comparative safety of the industries that they support. It can be argued that increasing numbers of incidents and accidents reflect inadequate regulation as well as unsafe working practices within an industry. Many regulators are sensitive to these criticisms. The UK Health and Safety Executive's report in Signals passed at Danger (SPADs) reveals the tensions that exist when regulators monitor and publish incident information [349]. The document begins by stressing the relative safety of rail travel and overall improvements in the frequency of these incidents. In 1997-98, Her Majesty's Rail Inspectors (HMRI) received 593 reports of SPAD incidents across the UK rail network. This represented a reduction from the high of 944 incidents in 1991-92. This generally positive tone is balanced by the following paragraphs of the report which acknowledge that 'while such incidents continue to occur, there remains the possibility of one leading to a train collision and/or derailment'. They also note that the number of SPAD incidents increased to 643 in 1998-99, reversing the earlier downward trend. This careful balancing of positive and negative statistics continues throughout the report. It continues by noting that many of these 643 incidents do not threaten safety because the train stops within the 183 meter overlap to the signal which is the intended safety margin. These more positive comments are then balanced by the observation that potentially severe SPADS involving trains that run past the overlap and where there are connections ahead increased from 42 in 1997/98 to 52 in 1998/99.

The previous paragraph illustrated the *regulators dilemma*. Falling incident statistics illustrate the effectiveness of a regulator and their reporting system. However, by focusing on these figures there is a danger that the regulator may appear complacent in the face of any subsequent accidents. Conversely, rising incident statistics can be interpreted as the result of ineffective regulation even though they may indicate that sufficient information has been obtained to guide subsequent intervention. This dilemma can lead some regulators to stress the difficulty of intervening to prevent particular types of adverse events. For instance, there were 230 collisions at Canada's 22,400 public highway and railway crossings in 1998. Another 46 collisions occurred at private and farm crossings of railway lines. The National Safety Program, Direction2006, stresses that more than 50% of these incidents occurred at crossings that were equipped with automated warning devices such as flashing lights, bells and gates. There were a further 80 collisions involving trains and pedestrians. It is concluded that 'the fact that so many vehicles and pedestrians are involved in collisions with trains while either disobeying railway crossing signs and signals or trespassing on railway property underscores the need for increased enforcement' [212]. This response to the regulator's dilemma shows how adverse incident statistics can be used to justify different forms of intervention, such as 'increased enforcement' when existing measures appear to have failed. It remains to be seen whether this particular approach will have the intended effect.

It is important to stress that regulators, like companies and safety managers, often have several objectives for monitoring incident reporting systems. As we have seen, incident data can be tracked to identify areas for intervention or to monitor progress towards particular safety objectives. It is also important to monitor the performance of reporting systems and not simply the data that they produce. Regulators must account for their expenditure in terms of the 'productivity' of their reporting systems. For example, the annual report of the Chief Executive for New Zealand's Transport Accident Investigation Commission focuses on these metrics [628]. This account opens with the observation that the Commission launched 47 investigations, finalised 36 reports and promulgated 112 safety recommendations for a total cost of $1.588 million in 2000-2001. This represented an overspend of 0.1% beyond the Commission's income of $1.586 million.

The ATSB provides a further example of meta-level monitoring in which the performance of the reporting system is analysed as well as the individual incidents [51]. The 2001 annual report identifies a number of specific metrics that are to be used in assessing ATSB activities in the following twelve months. For example, one core activity was identified as the investigation of rail safety incidents to ‘identify circumstances and establish causes’. The annual report identifies quality, quantity and timeliness metrics. Quality can be assessed by ensuring ‘impartial investigations undertaken in accordance with relevant legislation/regulations and procedural guidelines’. The quantity criteria are set as ‘Findings published in up to 4 reports’. Timeliness metrics establish a median time of 27 weeks to complete investigations and finalise reports. Another key activity was to ‘Facilitate and publish rail safety statistical analysis and data collection to assist in the conduct of rail safety investigation and the development of policy and strategies’. The quality of this activity was to be assessed in terms of ‘user satisfaction with published statistical information’. The quantity criteria was again established as publishing 4 statistical reports. An associated comment noted that in the previous year; ‘work on a rail safety statistical database development continued but delay in agreement with state rail accreditation authorities delayed publication of data’. Similarly, a further core activity was to ‘publish and distribute rail safety reports’. The quality of this activity was to be assessed in terms of the acceptance and utilisation of rail safety reports by the rail industry. The plans include a commitment to publish the findings in up to 4 reports.

Political Monitoring. Politicians, typically, help to establish the regulatory structures that protect public safety. They, therefore, have a keen interest to ensure that monitoring data reflects the success of those structures. When evidence is presented about particular short-comings then there is often a rapid move to ensure that appropriate action is taken. For example, John Spellar, the Minister for Transport, recently told a rail industry safety conference that government and industry must act to reduce the 300 deaths from trespass and suicide on UK railways each year. He argued that over half of these incidents were due to malicious acts of criminal damage and that it was, therefore, necessary to introduce a coherent ‘national strategy’ to address the problem [218]. The political sensitivity over incident data also partly explains regulators’ concerns to both justify their intervention and to account for their expenditure on reporting systems. Ultimately, accidents can lead the general public to question the political structures that guide the development of safety policies at a national and an international level. For example, the Indian Government of Atal Bihari Vajpayee ordered a complete review of their national rail system in the aftermath of the Gaisal train collision in which almost 300 people died in 1999. He refused to accept the initial offer of resignation from his railways minister, Nitish Kumar, who said ‘he felt the need to punish himself for the huge loss of life’ [101]. It is instructive to note that this political reaction was triggered in spite of a relatively good safety performance across the Indian rail network. In 1997, there were 1.4 passenger deaths for each billion passenger-kilometres travelled in India compared to 1.42 in the European Union.

The events surrounding UK rail privatisation provide a more complex example of the role that politics play in monitoring the performance of incident reporting systems. The break-up of British Rail, the national rail service, was proposed throughout the late 1980’s but only emerged as a commitment in the 1992 Conservative election manifesto. A White Paper on rail privatisation was then produced following their victory under John Major in July 1992. Pressure from the Treasury resulted in a decision to separate the operation of the infrastructure from that of rail services. This led to the creation of twenty-five separate companies, including Railtrack which assumed responsibility for the rail infrastructure. The Bill was finally passed in November 1993. The first operating franchises were offered in December 1995 to SouthWest trains, LTS and Great Western. Railtrack was floated in 1996. The final operating franchise was offered to ScotRail in April 1997 shortly before the Labour party was elected to power. There then followed a succession of high-profile failures including accidents at Watford Junction in August 1996, Southall in September 1997, Ladbroke Grove in October 1999, Hatfield, in October 2000 and at Selby in February 2001.

These ‘failures’ helped to launch a series of enquiries and investigations that considered the monitoring of incidents and accidents as part of a wider review of rail safety in the UK. Previous paragraphs have cited from the Cullen report into Ladbroke Grove and the Health and Safety Executive’s report into Signals Passed at Danger. Rather than reiterate the findings of these investigations, it is also important to consider the political impact of these initiatives to monitor both incident information and the reporting systems that produce them. Don Foster, the Liberal Democrat transport spokesman, reviewed these statistics during a Commons debate into transport safety. He concluded that the way in which the Conservative government had introduced privatisation had created 'confusion between safety and other aspects of the railway - not least confusion between safety and profit' [99]. The privatisation process had created uncertainty about who was responsible for what happened when an accident occurred. Labour's junior transport minister, Keith Hill, responded by arguing that both public and private transport operators must make safety their first priority; it is 'totally unacceptable for financial interests to take precedence over safety'. His response also illustrates the way in which narrow discussions about the safety record of particular companies can be broadened by political debate. He was compelled to defend plans for the 'part-privatisation' of the London Underground and for National Air Traffic Services (NATS) in a debate on the Ladbroke Grove rail crash. The political nature of such incidents is again illustrated by the Conservative spokesperson, Shaun Woodward, who argued that 'the public not only wants us to be concerned about safety but to ensure that when we know that safety may be at risk, to take responsibility and action when and where appropriate' [99].

The Hatfield accident, in particular, helped to focus attention on the high levels of investment that were necessary to achieve acceptable safety standards throughout the UK rail infrastructure. Political pressures ultimately forced the Labour government to withdraw financial support from Railtrack. The infrastructure company was then taken into administration. This political decision had both financial and operational implications. Over 250,000 shareholders, which included 90% of the company's employees, were immediately affected by this decision. The withdrawal of government financial support for Railtrack also cast considerable uncertainty over the future of the UK rail network. In the aftermath of this action, the percentage of trains arriving 5 or more minutes late increased from approximately 25% to over 30%. Although these figures were subsequently challenged on the grounds that they reflect a seasonal increase in delays from adverse weather conditions and 'leaves on the line'. The decision also raised safety concerns that demoralised employees facing an uncertain future might exacerbate existing equipment and infrastructure problems to trigger an increase in adverse events. Hence the political intervention directly led to a request from the Department of Transport, Local Government and the Regions to the Health and safety Executive to increase their monitoring of rail incident data to ensure that the administrative procedures had not jeopardised the safety of the rail system. It remains to be seen whether the incident data will reveal the same adverse trends that many have claimed for reliability statistics.

Political interest in the data that can be obtained from monitoring reporting systems does not just focus on the need to counter potential criticisms of particular initiatives. Statistical evidence of falling reporting rates is often used to validate previous actions. It can also be used to publicise and promote a reporting system. This may indirectly increase confidence in the wider regulatory systems that protect the public. For instance, in 1999 the U.S. Transportation Secretary and Federal Railroad Administrator announced the publication of a report showing 'dramatic' improvements in railroad safety as a result of the Clinton administration's partnership with industry. In 1997-98 there was a 27% reduction in railroad employee fatalities and a 33% reduction in passenger fatalities. Highway-rail incidents declined 9% and highway-rail injuries 15%. The FRA also reported a 'dramatic' fall in six-year safety results. From 1993 to 1998, highway-rail incidents declined 28%, highway-rail fatalities 31% and highway-rail injuries 29% while railroad operations, measured in train miles, increased 11%. It was argued that the Clinton partnerships supported safety improvements by focusing attention on the 'root causes of problems' and an improved understanding of 'the nature of rail-related incidents'. The Transportation Secretary stated that:

> "'President Clinton and Vice President Gore challenged the government to do business in a new way, to work better together and get results that Americans care about. The report we are issuing today demonstrates that this approach to governing is working by dramatically increasing safety in the railroad industry." [238]

As we shall see, however, headline figures can mask other incident statistics that often contradict political claims about the safety of an industry. Closer inspection of the FRA monitoring data shows that the overall fall in accidents and incidents was largely accounted for by the drop in highway-railway incidents from 3,865 (1997) to 3,508 (1998) The same period saw an increase in train accidents from 2,397 (1997) to 2,575 (1998) mainly caused by derailments, 1,741 (1997) and 1,757 (1998), and human factors, 855 (1997) and 971 (1998) [243].

Media and Public Involvement. There are clear reasons why those who are involved in the operation and management of a reporting system should want to monitor both its output and performance. It is also important to recognise that there may be other parties, including trade associations and public pressure groups, who have an interest in tracking this information. Many of these groups have indirect access to incident information. The US Freedom of Information Act has helped to ensure that many Federal agencies provide incident information over the web. The provisions of this act have had numerous benefits. For example, much of the recent research on novel computational techniques for incident retrieval has been directly driven by these new information sources [413]. It would not have been possible to write this book ten or even five years ago when there was little or no access to such confidential databases. Even where direct access is denied, pressure groups can monitor reporting systems indirectly through official press releases and less authoritative leaks to the media. For example, the BBC reported that ScotRail were one of '10 train companies warned by the Railways Inspectorate that it was not doing enough to prevent drivers passing red lights' [112]. HMRI's figures showed 56 SPADs in May 2001 compared to only 35 in May 2000 and an average of 49 SPADs between 1995-2001. The Railway Inspectorate warned operators that they would face enforcement actions and prosecutions if their safety records did not improve. The report goes on to explain that these criticisms were triggered because the number of SPADs had *improved* but only slightly. Media organisations do not always follow the balanced approach illustrated by this example. It is also important to recognise that concerns over this publicity encouraged ScotRail to directly counter criticisms in the BBC report. The following quotation presents the response of a ScotRail spokesman to the publication of the HMRI figures. The confidential reporting system is the CIRAS scheme that has been mentioned in previous chapters and will be discussed in later sections of this chapter:

> 'I think the figures they have surround the long term average rather than last year's results. Last year we had a 22% reduction, which was better than the national average. It is a subject that is taken very seriously. It is obviously very high up our agenda. We put a great deal of effort into it and we have led in the past on man y new initiatives, including the confidential reporting system.' [112].

These comments elicited a sympathetic response from passenger 'pressure' groups. The Deputy Secretary of the Rail Passengers Committee for Scotland acknowledged that the number of SPAD incidents had fallen since rail privatisation. He also referred to initiatives by companies, such as ScotRail's defensive driver techniques, that had helped to reduce these adverse events.

This example illustrates the diverse groups that are concerned to monitor data from incident reporting systems. The SPAD frequency information was initially released by HMRI. This investigatory and regulatory organisation is primarily responsible for controlling the hazards that affect the health and safety of anyone who might be affected by the operation of Britain's railways. The BBC then identified the information as having a wide public interest. This media organisation then commissioned a report which elicited responses from the companies concerned. They countered the HMRI's interpretation of the statistics by pointing to longer term trends. Finally, a passenger group responded to ScotRail's defence of their safety record. Regulators, investigatory organisations, the media, commercial organisation and public pressure groups all contributed to the analysis of information that was initially obtained from the SPAD reporting system. Such diverse opinions illustrate the difficulty of interpreting such statistics. Many of these problems stem from the paradoxes of incident monitoring.

Paradoxes of Incident Monitoring

It can be argued that the monitoring of incident reporting systems should ensure that they help to avoid future incidents and accidents. Chapter 1 has, however, argued that we cannot achieve absolute safety [675]. It is also important to emphasise that incident reporting systems do not operate in isolation from the rest of an organisation. A new scheme might be introduced at the same time as new processes and plant come on-line. Hence, the introduction of the reporting system may coincide with a notable increase in adverse events. **First paradox of incident monitoring:**

even if a reporting system does not demonstrate a long term reduction in adverse events it can still be argued that the safety record would have been even worse if the reporting system had not been in place.

Given that we cannot achieve absolute safety, it is important to shown that the level of investment in a reporting system yields an optimal reduction in adverse events. In other words, we would like to demonstrate that additional investment would provide little additional safety information. Conversely, we might also demonstrate that savings could not be made without jeopardising important feedback about the safety of the system. Unfortunately, a number of problems complicate the task of assessing the marginal utility of investments in incident reporting systems. In particular, there is an important distinction between the numbers of incidents that occur and the numbers of submissions made to a reporting system. Chapter 5 has described how increased levels of funding typically elicit additional submission. **Second paradox of incident monitoring:** additional funding for incident reporting systems typically yields an increasing number of submissions as people become more aware of the system. This need not reflect a rise in the underlying number of incidents. Conversely, cuts in the funding associated with a reporting system may yield fewer contributions but this need not indicate an improvement in the underlying safety of an application. Staff may be disillusioned with the effectiveness of the reporting system.

The monitoring of incident reporting systems is further complicated by the argument that in any resource limited environment, we cannot simply consider the costs of any particular activity in isolation. In contrast, it is important to assess the opportunity cost associated with maintaining an incident reporting system. This focuses on those activities that must be sacrificed in order to support a reporting scheme. I would stress the importance of this perspective given that the individuals who help to establish and maintain reporting systems are often amongst the most highly-trained and safety conscious staff within an organisation. These individuals are often so committed to the operation of a scheme that few seem to consider whether their time and energy might not be more effectively employed in other safety-reelated tasks. **Third paradox of incident monitoring:** those individuals who are most committed to the operation and maintenance of a reporting system may not be in the best position to judge whether or not these schemes make the most use of their finite resources.

The previous paragraphs have focussed narrowly on safety improvements as the principle benefit of operating an incident reporting system. As we have seen, however, there are many other reasons why one of these schemes might be established. Regulators might require operators to support a reporting system. Reporting systems can be introduced to deflect criticism of previous safety related failures. These schemes can also be introduced to form part of a wider 'lessons learned' or quality assurance scheme. In such circumstances, safety benefits form part of wider improvements in operating practices. **Fourth paradox of incident monitoring:** incident reporting systems may continue to be maintained even though almost no safety-related contributions are submitted. The rationale for operating the system need not rely narrowly upon safety-related issues but may have more to do with wider operational and regulatory concerns.

These paradoxes make it difficult to interpret the results of any attempts to monitor the success or failure of a reporting system. For instance, a fall in the number of incidents reported might indicate disillusionment with the system, problems in submitting report forms or a genuine reduction in safety-related incidents. The difficulty of interpreting particular measures has led some organisations to adopt a broader perspective. In particular, they have sought metrics that might be used to validate the diverse range of proposed benefits from incident reporting that were enumerated in Chapter 2. For example, the following list extends the results of a study by the US Coast Guard [834] to identify ways of monitoring the health of their reporting systems:

- *Number of submissions received.* This is often the most convenient means of monitoring participation in a reporting system. As we have seen, however, it can be difficult to interpret the results. Low submission rates may indicate safety improvements or disinterest in the system. Similarly, increases in participation may stem from the expansion of an industry as more groups are exposed to potential hazards. There are further dangers. For instance, long-running schemes often publicise their success by reiterating the cumulative total of reports received. If one looks more closely into the nature of the reports received, it is often depressing to find that the same sorts of failures have been submitted often over decades [409]. Hence a high cumulative total and high annual participation rates may indicate the limitations of the approach rather any measure of success.

- *Change in the quality of submissions.* Rather than focusing on the raw numbers of submissions, the success of a reporting system can be assessed in terms of the quality of those submissions. This can provide feedback on whether or not potential participants can understand and follow reporting procedures. At first sight, it may be argued that such measures provides relatively little information about the safety of an underlying application. Given the problems associated with measuring the frequency of 'near misses' this approach can, however, help to minimise any potential barriers that might otherwise prevent individuals from submitting information about such events. In consequence, it can be argued that if the quality of submissions improves then we can have greater confidence in the accuracy of reporting frequencies. A high number of apparently spurious submissions might lead to some potentially valid incidents being discarded during any initial filtering. If a contributor fails to provide sufficient information about a potential incident then analysts may be forced to invest scarce resources in collecting sufficient initial information to justify subsequent investigation. In many systems, the decision may be made not to invest those resources so that additional attention can be paid to more 'clear-cut' incidents. Similarly, a high number of spurious submissions might indicate that some 'valid' incidents are not being reported because of general confusion about the purpose of the scheme.

- *Percentage of attributable reports in an anonymised system.* In systems that offer contributors the possibility of filing a report without disclosing their identity, the proportion of submissions that include contact information can provide a measure of confidence in the system. This measure can provide indirect insights into participation levels. A high proportion of unattributable reports might indicate general skepticism about the integrity and potential benefits of the system. These concerns are likely to jeopardise participation in the system. It can, therefore, be argued that the number of incidents being reported to the system is unlikely to provide an accurate impression of the total number of adverse events and 'near misses'. Conversely, a high level of attributable submissions may indicate high levels of participation. This, in turn, can increase confidence that 'near miss' incidents are being submitted and that contributions provide a more accurate measure of underlying safety.

- *Number of submissions investigated.* Chapter 10 has argued that participation in many reporting systems depends upon organisations acting on the information that they receive. In this view, the effectiveness of any reporting system cannot be measured simply by the number of submissions that are made. Participation levels are unlikely to be sustained if no actions are taken to investigate the safety concerns that are identified by contributors. This view is significant because it emphasises the idea that the 'health' of a reporting system will change over time. The future success of the system may, therefore, depend partly on current submission rates and partly on the way in which the system reponds to those contributions.

- *Number of reports leading to safety improvements.* The caveat that there must be some demonstrable recommendation or action taken in response to a report is significant because otherwise a rise in submissions might reflect an increase in spurious reports. The numbers of reports that trigger interventions not only provides a measure of the effectiveness of any system, this information can also be used to encourage further participation. Such information can demonstrate that reports will be acted upon. The 'safety improvements' that are derived from a reporting system can be interpreted quite broadly. For instance, the UK Rail Safety group monitors the number of enforcement actions that HMRI initiates against operating and infrastructure companies. 30 notices were issued in the second quarter of 2001/02 bring the six month total to 43 notices. This can be compared to only 31 notices for the whole of 2000-2001 [691]. It can, however, be difficult to draw firm conclusions from these figures. Not all enforcement actions are triggered by incident reports. Conversely, not all incident reports that identify potential violations will lead to enforcement actions. Similarly, a rise in the number of enforcement actions can be the result of short term initiatives by investigatory agencies rather than the result of short-term increases in the number of adverse events. For instance, the increase in enforcement actions in 2001-2002 was partly the result of actions to reduce the frequency of trespass and vandalism.

- *The number of reports submitted by particular categories of participants.* For example, the success of a reporting system might be measured for particular industry segments, regional areas, professions or staff positions. Such measures are important because most successful reporting systems achieve safety improvements in spite of 'uneven' levels of participation. For example, the Aviation Safety Reporting System gathers very few reports from Military pilots and a limited number from General Aviation. The FDA's MEDWATCH program receives proportionately less reports from nursing homes than it does from larger hospitals. In rail reporting systems, there are few reports of 'Signals Passed At Danger' in remote regions where there are few witnesses to any infringement. In such circumstances, the health of a reporting system may be judged against participation targets for particular groups of participant.

- *Change in the number of accidents.* As mentioned previously, changes in the number of submissions to a reporting system can be the result of other events that have little to do with the underlying safety of any application. It can also be difficult to gather accurate statistics about the occurrence rates for 'near miss' incidents. In consequence, the only reliable safety measure is the number of accidents within an industry. For instance, the Cullen report argued that several major accidents indicated significant flaws in existing reporting practices within the UK rail industry. As we shall see, however, structural changes in this industry created new hazards and placed new demands on the existing reporting infrastructure. This emphasises the importance of continually monitoring the performance of a reporting system against such measures. A reporting system may provide adequate information about potential hazards within one context of operation but may be ineffective in identifying potential hazards as changes occur within an industry. Further problems arise because the low frequency of accidents in many industries can prevent reliable inferences being made about the underlying safety of an application until an adverse event occurs. It can also be difficult to define what constitutes an accident. Some injuries and fatalities, for instance from suicide or trespass, are difficult for operating companies to control. The practical problems of calculating an accident rate as a means of assessing the performance of incident reporting systems can be illustrated by the UK Railway Safety Group's quarterly reviews [691]. This calculates the risk of a train accident for the previous quarter by combining the frequencies of particular contributory factors, including level crossing mis-use, irregular working activity and vandalism. The complexity of using such measures to assess overall 'safety' is illustrated by the October 2001 report. This recorded a slight increase in the accident risk even though the number of 'significant train accidents' actually fell. This apparent paradox can be explained by a rise in workforce fatalities. There are further complications. The incidence of track quality faults, wrong-side signal failures and train speeding reduced fell but the number of public accidental fatalities rose compared to in the first quarter. The practical difficulties in compiling accident statistics are exacerbated by ethical objections. Arguably the most significant criticism of accident metrics is that the performance of a reporting system is assessed in terms of the number of times it fails to protect either the workforce or the general population.

- *Change in the number of particular event types.* It may not be possible to gain an accurate assessment of the overall number of 'near miss' incidents and accidents across an industry. The problems of under-reporting and reporting bias frustrate attempts to gather such statistics. These general problems can be addressed by focusing on particular types of adverse event. Additional publicity can be provided to explain the importance of certain hazards. Automated monitoring and logging systems can be used to detect when such events have occurred. The results of these special initiatives can be compared against levels of participation to provide a measure of any previous under-reporting. Unfortunately, the effectiveness of these techniques may decline if they are used too frequently [409]. Participants may becomes immune to successive attempts to sensitise them towards particular types of failure.

- *Changes in the safety issues identified.* A danger with any reporting system is that it will continue to identify the same safety concerns that have been observed in previous incident reports. If participants continue to reiterate well known issues then it might be argued that the reporting system is ineffective as a means of addressing those issues. This view can be challenged. For instance, there may be agreement over the nature of the problem but disagreement over the recommendations proposed by incident investigators. For instance, Chapter 9 described how the National Transportation Safety Board (NTSB) struggled to introduce devices that were intended to address gas leaks and explosions that occurred over almost three decades. Industry representatives argued that the costs associated with such changes would not be justified by any potential benefits. It can, therefore, be argued that the continuing pattern of incidents did not simply reflect the failure of the reporting systems. Instead, it indicated the difficulty of resolving complex commercial and regulatory issues and the need to build up a body of evidence in support of the investigators' recommendations.

- *Changes in outcomes.* The success of a reporting system might be measured at a gross level in terms of a reduction of the total working days lost to industrial injuries. Similarly, it might be measured in terms of any change in particular types of injury or fatality. For instance, Table 15.2 presents five-year trend data on rail fatalities and serious injuries in Australia. These are categorised according to individual regions. Unfortunately, it can be difficult to obtain such outcome information. Some of the values in Table 15.2 denoted by the periods have been suppressed because of State privacy restraints. Although this data was published as part of a national report on transportation safety, the statistics had to be pieced together from the Australian Bureau of Statistics and the Australian Institute of Health and Welfare. These were the only sources of national rail safety data available in the absence of a national rail occurrence database.

Fatalities

	NSW	Vic.	Qld	SA	WA	Tasmania	NT	ACT	Aust.
July 1993 - June 1994	10	8	10	4	5	0	0	0	37
July 1997 - June 1998	22	12	2	1	5	0	0	0	42

Serious injuries

	NSW	Vic.	Qld	SA	WA	Tasmania	NT	ACT	Aust.
July 1993 - June 1994	80	22	24	7	10	2	0	0	145
July 1997 - June 1998	66	18	19	..	13	3	0	2	..

Table 15.2: Rail Incident Outcomes on Australian Railways (1994-1998)

Outcome measures can also be used to assess the performance of incident reporting systems in individual companies. A rising number of serious injuries might be interpreted as a failure to learn from previous incidents. Again, however, there is considerable concern over the reliability of this approach [338]. For example, it can be argued that the outcomes might have been even worse if the reporting system had not been in place. It can also be difficult to identify suitable outcome measures that might be used to assess the performance of individual firms. The outcome of an adverse event can be mitigated by the prompt intervention of medical staff. Conversely, the eventual outcome of some incidents may take many years to fully develop. There are some industries, in particular those that depend on self-employment labour, for which it has always been difficult to obtain accurate consequence statistics. Further concerns stem from the difficult of accounting for near misses with high-potential consequences. For example, no-one was killed or fatally injured by a main track derailment on Canadian railways between 1983 and 1996. During that time, there were approximately 10 derailments per year from bearing failure alone [776]. It can also be difficult to distinguish the impact of a reporting system on any changes in consequence figures. It is for this reason that the Transport Canada requires rail operators to monitor the performance of their reporting systems, in terms of the types of failure and remedial actions, as well as employee deaths, disabling injuries and minor injuries per 200,000 hours worked [779].

- *Survey results from user groups.* Given the problems associated with deriving accurate measures from either the submission rate to a reporting system and the ethical issues associated with post-hoc accident rates, it is important to find other means of assessing the effectiveness of these initiatives. Given that many reporting systems have identified failures in ill-defined concepts such as 'safety culture', it can be argued that such schemes are successful if they act to change those previous weaknesses. These schemes remind participants of previous incidents within their industry and hence can play a positive role in informing people about the potential adverse consequences of particular incidents. This 'consciousness raising' effect can be assessed by surveys of the groups who participate in a reporting system. This approach recognises that a far larger group may benefit from the publications produced by reporting system that the comparatively small number of individuals who might actually witness an adverse event and then submit a report.

- *Change in insurance premiums.* The previous measures focus on attributes of reporting systems or on the applications that they are intended to protect. It can be difficult to gather accurate figures for these direct measures. It can also be difficult to interpret what changes in these measures imply for the safety of an application. Some organisations, therefore, emphasise the indirect benefits of incident reporting systems. These include reductions in insurance premiums associated with safety-critical applications. Such 'metrics' are credible because they typically reflect the judgement of an external organisation that is strongly motivated to provide an accurate risk assessment.

- *Changes in application operating costs.* Incident reporting systems are often integrated into more general systems for quality control. Improvements in operating efficiency are often more easily measured that any improvements in safety. For example, the relatively low frequency of many safety-related events can imply that individual units will only receive a few submissions each year. It is, therefore, impossible to judge the relative success of a system from month to month. In such circumstances, organisations are often motivated to increase the scope of a 'lessons learned' system. It is hoped that by reporting lower consequence failures, potential participants will be more comfortable with the procedures that support the submission of safety-related events. It follows that even if no safety information is submitted to the system, the provision of information about other potential problems in quality control or efficiency can provide feedback about the effectiveness of the reporting system.

- *Change in the operating cost of the reporting system.* Incident reporting systems do not operate in a commercial vacuum. In consequence, many systems are assessed according to the usual financial criteria associated with any management or engineering function. Unfortunately, the problems in deriving objective measures for the success of a reporting system can make these schemes vulnerable to cost cutting. It can be difficult for safety managers to prove that cuts in the funding of a reporting scheme will jeopardise the safety of application processes. Similarly, a large increase in the number of reports processed at the same level of funding raises questions about the level of analysis that can be sustained for any particular safety issue. Such savings can be justified through increased efficiency, for instance by the introduction of information technology. Alternatively, more accurate forms of risk assessment can be used to ensure that reduced funding does not impair the organisation's response to high-criticality incidents.

- *Establishment of a self-sustaining operation.* The success of some reporting systems is measured against particular commercial or financial criteria. Increasingly, there is a view that these systems should be self-sustaining and should not be sustained by public money. The industries that benefit from the insights obtained by a reporting system should meet the costs associated with maintaining the system. This creates potential concerns. For instance, if some companies 'opt out' of the scheme then they may be isolated from any insights provided by the system. If companies are forced by the regulator to join the scheme then this can be interpreted as undue interference in the commercial operation of particular industries, especially if the regulator retains an interest in the maintenance of the scheme. Conversely, if a commercial cartel retains control of the reporting system there is a danger that the independence system can be compromised. Investigators may be unwilling to propose recommendations that have high cost implications for the rest of the industry.

- *Changes in the mode of submission.* Given the difficulties of obtaining and interpreting objective measures for any safety improvements derived from a reporting system, it is often more convenient to identify more focused objectives that relate to the way in which a particular scheme is implemented. This class of measures are often associated with the financial objectives, summarised above. For instance, the success of a reporting system can be assessed in terms of particular modes of submission. Several of the schemes described in this book have moved away from paper based submission towards telephone, fax and Internet based contributions [423]. These initiatives are intended to increase the scope of a system by cutting the costs associated with managing the collation of individual reports. The use of these metrics indicates the complexity of monitoring incident reporting systems. These changes can introduce new biases into the reporting process, it may be harder for some participants to access and use new submission techniques. By achieving particular objectives for the introduction of new technology, the reporting system may lose important safety-related information. This may, however, only be a short-term effect as more people learn how to operate the revised submission procedures. The reduced costs associated with alternative modes of submission may be necessary to support the long-term survival of the system.

- *Number of information requests.* Previous measures have focussed on the number of incidents reported or the number of investigations that have been completed. The success of a reporting system can also be assessed in terms of the information that it disseminates. Gathering information about previous failures is of little benefit if any insights are not passed on to those who are best placed to use them. For this reason, the success of a reporting system might be measured in terms of the number of information requests that are received. As with the submission metrics, more fine-grained targets might also be associated for requests from particular end-user groups within particular industries or regions.

- *Time to implement changes from first notification.* A number of temporal properties of incident reporting systems can be measured to provide insights into their efficiency in dealing with particular safety-related concerns. For example, it is possible to record the time between a request being made for incident information and that request being addressed. Such intervals are significant because any delay might compromise the safety of application processes. Similarly, the time between an initial notification and any secondary investigation could be measured to provide information about the response to a report. This is a significant concern given that safety managers have found completed report forms that have lain neglected for many months in the desks of process supervisors. These metrics, typically, introduce additional administrative overheads in terms of the resources that are required to log timing information. They are, therefore, most frequently gathered by large, distributed systems such as national Air Traffic Management reporting schemes [423].

- *Number of publications issued and acted upon.* If reporting systems disseminate most of their information through paper-based publications it can be difficult to gain a true measure of all of the individuals and organisation who may read and act upon the information that is disseminated. Each journal or bulletin can be read by several people. Conversely, there is no guarantee that the recipients of a publication from a reporting system will have actually read the information that it contains. Readership surveys provide one means of addressing these problems. Alternatively, the 'productivity' of a reporting system might be measured in terms of the raw number of publications that it produces. Unfortunately, such measures do not discriminate between active systems that continually provide new insights and those that regularly publish the same advice without seeking new remedies for past and present failures. These publication measures might be supplemented by an assessment of the other 'peripheral' activities that often provide alternative means of dissemination. Conference presentations and workshops can provide a further indication of the health of a reporting system.

- *Number of people who access computer-based resources.* It can be difficult to track all of the people who have access to the paper-based publications that are produced by an incident reporting system. Some of these problems can be addressed through the provision of computer-based resources that automatically log any requests for information. The metrics provided by these systems can help to justify any investment in computer-based resources. In particular, it is important to demonstrate that the use of electronic dissemination techniques does not hinder access to information about previous incidents. Automated logging facilities can be used to provide profile information based on the Internet Protocol address of sites that request access to the system. These addresses uniquely identify the computer that sent the request. More accurately, they identify a connection between that computer and the network because a single machine might have several network connections. In practice, however, the allocation of IP addresses to sites and the local routing of requests can limit the inferences that are made. Automated logging can provide other metrics. For instance, it is possible to identify the number of abandoned or failed requests made to a web server. This provides useful information about retrieval delays. If the number of abandoned requests rises then it may be necessary to index the data in another way or to provide additional support for the system infrastructure.

- *Changes in industry/operator participation.* Many of the proponents of incident reporting have argued that active participation from industry is required in order for these systems to be successful [844]. The imposition of reporting systems by regulatory intervention can lead to resentment and the creation of informal barriers that may discourage submissions from some employees. In contrast, the enthusiastic promotion of a reporting system can encourage participation and support the dissemination of safety-related information. In consequence, many regulators publish lists of the companies that have chosen to 'sign up' to a scheme. These lists provide a gross indication of industry participation. They provide little indication of the financial and organisational resources that each company is prepared to allocate to a reporting system. For instance, many hospitals have established incident reporting systems as a means of combating negligence claims. Many of these institutions provide limited budgets and appoint relatively junior staff to manage these schemes. In contrast, some hospitals have ensured that clinical risk managers are promoted to the highest levels within their organisational structure. Such differences make it difficult to derive accurate measures for the level of participation in incident reporting systems.

- *Levels of information sharing between companies.* Incident reporting systems have often been established with the claim that they will improve the dissemination of safety-related information between the participants in the scheme. It is, therefore, appropriate to consider how such information exchanges might be measured as a means of validating these claimed benefits. Chapter 5 has described how the creation of such systems can only have a limited effect on the barriers that prevent the effective dissemination of safety information. In consequence, many companies will operate their own internal schemes in parallel with industry-wide systems. This tends to ensure that only some incidents are shared in the manner proposed by the proponents of incident reporting systems. It would be very revealing to measure the differences between those incidents that are retained within a proprietary system and those that are shared in an industry wide scheme. Such initiatives would have to address the same barriers that prevent the exchange of information in the first place.

- *'Collateral' effects on industry.* Direct measures can be found for the impact that reporting systems have upon particular industries. Large numbers of similar incidents can trigger external regulatory intervention to enforce the recommendations that are made within an incident reporting system. Some health and safety organisations judge their success in terms of the numbers of prosecutions that are initiated in response to reports of adverse events. The information that is collected about 'near miss' incidents is often cited in legislative changes. In extreme cases, such reports can motivate government intervention to restructure an entire industry. The reorganisation of the UK rail infrastructure provides an example of such intervention [194]. Many of these changes cannot be initiated from within the incident reporting system itself, the influence of such schemes therefore extends well beyond those who are directly involved in operating the system.

- *Tracking of public image.* Lough has recently argued that the success of any reporting system should be measured in terms of its acceptability both by those who participate in the system and by the wider community in which it operates; "acceptability is often a reflection of high validity" [501]. His use of the term 'communittee' is interesting because it can refer to a 'communittee of practice'. His work focuses on techniques to support incident investigation by medical doctors in general practice. The term might also refer to the wider 'communittee' which includes the general public. This ambiguity is important because it identifies a dual role for incident reporting systems. On one level they can be used to derive particular insights that may prevent the recurrence of safety-related incidents. At another level, these systems act as an important means of reassuring the public that application processes are being operated in a responsible manner. This may, in part, explain why so many incident reporting systems have been established in the aftermath of major accidents. Such high-profile failures affect public confidence. Incident reporting systems satisfy their expectation that government and regulators should do something to address their safety concerns.

- *Longevity and 'technology transfer'.* The ultimate success of a reporting system can be measured in terms of its longevity. For instance, the US Aviation Safety Reporting System has continued in operation since 1976. The fact that it has survived through many changes in the fortunes both of the aviation industry and its sponsoring organisations demonstrates the perceived success of this system. The Australian Incident Monitoring System (AIMS) has a similar 'track record' within the field of patient safety, stemming from an anaesthesia project in 1989. Both of these applications have provided templates for subsequent systems. For instance, the operators ASRS argue that the success for their system led to the UK's Confidential Human Incident Reporting Program (1982), the Canadian SECURITAS system (1995), the Australian Confidential Aviation Incident Reporting system 1988, the Russian Voluntary Aviation Reporting System (1992), the Taiwan Confidential Aviation Reporting Enterprise (2000) and the Korean Confidential Aviation Incident Reporting System (2000) [60]. Similarly, the Australian Pateint Safety Foundation (APSF) which helps to administer the AIMS application has inspired the UK National Patient Safety Agency (2001) and the US National Patient Safety Foundation
indexNPSF (1998) both of which are closely involved in medical incident reporting. Imitation might provide the greatest evidence for the success of particular reporting systems.

Previous paragraphs summarise the vast range of metrics that have been proposed to support the monitoring of incident reporting schemes. Unfortunately, the strengths and weaknesses of these various measures have not been established. For example, there is no evidence to support criticisms against raw submission numbers as an indicator of the contribution to system safety. This lack of evidence is unsurprising. The relatively low frequency of accidents prevents analysts from forming the causal connections that might support statistical correlations. We might like to establish that organisations with a low number of submissions also suffer from a higher frequency of more serious accidents. Things are not so straightforward. For instance, several of the UK's rail operating companies with the best reporting record have also experiences significant safety-related problems [417]. The difficulty of establishing measures that relate incident reporting behaviour to accident frequencies is further illustrated by Wright's [874] recent work on the Heinrich ratio, summarised in Chapter 2. She argues that railway workers are, typically, either involved in fatalities or are witnesses to 'near-misses' [874]. There are few reports in the middle ground of more serious, non-fatal incidents. If her analysis is correct then we cannot expect there to be any clear-cut relationship between submissions and accident frequencies.

It is difficult to conduct controlled experiments in this area. Several of the metrics proposed in the previous list can be influenced by local effects. For example, the support that an organisation provides for participation in a system can be affected by the behaviour of individual managers. One could envisage a trial which compared the influence that different supervisors had upon the reporting behaviour of their workforce. It is difficult to see how such influences could be distinguished from the mass of other local factors that might also affect reporting behaviour. These include the composition of work groups as well as the submission and reporting processes that operate in individual plants. Ethical problems also complicate work in this area. If participants are informed that they are being studied then this may affect their participation in the system. Conversely, post hoc studies can compromise the confidentiality of the reporting system if they associate particular reports with particular working groups.

The lack of direct evidence to support particular metrics reflects the wider lack of research to support many other aspects of incident reporting. Considerable resources have been devoted to support the design of safety-critical systems. Far less resources have been allocated to understand why these systems fail. In consequence, the development of incident reporting systems resembles a craft skill rather than an engineering discipline. Techniques are borrowed from other systems that are perceived to be successful. Often metrics are chosen because they either validate the allocation of resources to maintain the system or because they have been used to assess other similar systems. In many cases, there is also an unquestioning assumption that incident reporting systems are ‘a good thing’ hence it is largely irrelevant to look for more quantitative forms of support.

With these comments in mind, the following pages focus on a number of the measures proposed in this opening section. The analysis is grouped into three parts. The following section looks in more detail at the reasons why it is important to monitor the outcomes of incident reporting. In particular, we focus on the role that these systems play in risk assessment, in systems development, in training and in operational efficiency. The subsequent section justifies attempts to monitor the reporting process itself. Particular attention is paid to changes in submission rates, to the behaviour of investigators and to the implementation of proposed changes. The closing sections of this chapter present a range of techniques that can be used to implement the metrics that are identified in the previous sections. These range from the use of computer-based audits to monitor the behaviour of incident investigators through to observational studies of the working groups that submit incident reports in the first place.

15.1 Outcome Measures

Heinrich’s pioneering work in the area of safety management identified a number of tasks that safety managers must perform if incident and accident data is to inform the future operation of application processes [340].

1. collect incident and accident data;
2. analyse the data;
3. select appropriate remedies;
4. implement those remedies;
5. evaluate effectiveness of any remedies.

This approach can be criticised because it does not explicitly ‘close the loop’. In other words, it is implicit that the evaluation of any remedies will help to inform the selection of future interventions. Similarly, the evaluation process might itself help to inform or direct the elicitation of incident data. Kjellen addresses some of these limitations when he argues that the monitoring of an incident reporting systems must help to identify the need for further information as well as identify priorities for intervention [444]. This iterative approach suggests means of monitoring the effectiveness of an incident reporting system. Evidence collected during the first stage of Heinrich’s model can be used to provide insights into the effectiveness of previous interventions. As we have seen, however, it can be difficult to rely solely on changes in the numbers of submissions that are made to reporting systems. Contribution rates can change independently of the underlying number of safety-related incidents. The development of a reporting system can increase staff awareness of the need to report particular types of failure.

There are several alternative outcome measures that can be used to evaluate the effectiveness of incident reporting systems. Indirect observations provide feedback about those factors that have contributed to previous incidents and accidents. For example, attitudinal surveys and proficiency tests can be used to determine whether staff are better equipped to deal with situations that led to past failures. Similarly, maintenance activities can be monitored to determine whether they offer effective protection against previous incidents.

A limitation with the use of indirect measures is that previous incidents seldom recur in precisely the same way [699]. It is, therefore, important because revised training and operating procedures cannot simply be based on previous incidents, they must also consider alternative failure scenarios. This implies that incident reporting systems should not only be assessed in terms of the feedback that they provide about existing operations, they should also be evaluated in terms of the contribution that they make to feed-forward risk assessment. This has recently led to the development of accident prediction tools, such as the FRA's Highway-Rail Crossing Web Accident Prediction System (WBAPS) [244]. This uses historic data about previous incidents at particular types of rail crossings to anticipate future accidents at similar locations. Such applications raise ethical questions that complicate the monitoring of incident reporting systems. If an accident prediction proves to be correct then the overall regulatory system can be criticised for failing to prevent a failure that had been anticipated. Ideally, such incident data should direct acquisitions and design policy so that such 'anticipated accidents' are avoided. In particular, reporting systems should inform risk assessments so that the weaknesses of previous systems are not replicated in future developments. A further form of indirect monitoring is, therefore, to assess the impact that incident information has upon future systems and not simply the operation of existing applications.

15.1.1 Direct Feedback: Incident and Reporting Rates

Many industry regulators publish annual summaries that can, in part, be used to assess the performance of reporting systems. For instance, Table 15.3 provides an overview of the US Federal Railroad Administration accident and incident data for 1999 and 2000 [243]. These statistics illustrate some of the problems that arise in interpreting 'raw' information about failure rates. There was a reduction in the total number of reported casualties, from 12,632 to 12,580. At the same time, however, there was an increase in the total number of fatalities, from 932 in 1999 to 937 in 2000. It might be argued that these figures represent an improvement in the safety performance of the rail industry, as noted by the 0.5% reduction in casualties mentioned in Table 15.3. Alternatively, it can be argued that the rise in the number of fatalities represents a worsening of the overall safety record. The reduction in the total number of reported casualties, in this more negative interpretation, might reflect a reluctance to report important safety information.

The relatively small changes illustrated by these statistics can also be explained by annual fluctuations in the incident statistics rather than by changes in the underlying systems. For example, Figure 15.3 normalises accident rates against the number of miles that were travelled in 1999 and 2000. It does not, however, account for differences in the time that it took to travel those distances. Small changes in the average speed of a journey can affect the risk exposure of both staff and passengers. This may be determined by changes in the weather from one year to the next. It might, therefore, be concluded that the small fall in reported casualties might be accounted for by such factors rather than by any overall improvement in rail safety. In order to guard against such apparently 'random' effects we must also consider the issue of statistical significance. We can identify two possible dangers in the interpretation of incident statistics such as those presented in Table 15.3 [373]. A *type 1* error occurs when we decide that changes in the operation of a safety system had an effect on the overall incident data when they did not. A *type 2* error occurs when we decide that changes in the operation of a safety system had no effect on the overall incident data when they did. Significance levels provide a measure of the probability of making a type 1 error.

Data:	Jan-Dec 1999	Jan-Dec 2000	%age Change
Train accidents	2,768	2,983	7.8%
Train accidents per million train miles	3.89	4.13	6.2%
Total reported casualties	12,632 (932 fatal)	12,580 (937 fatal)	-0.5%
Trespasser fatalities	479	463	-3.3%
Employee casualties per 200,000 employee hours	3.39	3.44	1.4%
Highway-rail crossing incidents	3,489	3,502	0.4%
Highway-rail crossing fatalities	402	425	5.7%
Highway-rail crossing incidents per million train miles	4.90	4.84	-1.1%

Table 15.3: FRA Accident/Incident Statistics, February 2002

Given the fluctuations that one might expect in the contribution rate for incidents and accidents, we might therefore set stringent requirements to avoid type 1 errors. There is, however, a trade-off. The lower we set the significance threshold for type 1 errors, the greater the chance there is of making a type 2 error. Further problems affect the use of such statistical techniques. Significance levels are most easily established for carefully controlled experimental situations in which it is possible to distinguish the change, or independent variable, that is linked to any measure, the dependent variable. Unfortunately, there are likely to be many factors that have an impact on overall incident and accident rates. For example, the UK rail sector has recently gone through profound structural changes. It is, arguably, impossible to distinguish the impact of these changes from other changes, such as the introduction of the CIRAS reporting system mentioned in previous Chapters. Assuming that we witness a reduction in the number of rail incidents in the UK, how can we determine whether that improvement was due to the introduction of the reporting system or to higher level changes in the regulatory environment? The problems of obtaining and interpreting incident and accident statistics affect most of the monitoring techniques that will be described in this chapter. For now it is sufficient to observe that these problems are currently being addressed by several recent initiatives. For instance, the statistical unit within the UK Health and Safety Executive has promoted the development of professional standards for the publication of safety-related information by both public and private organisations [338].

Industry-wide incident rates are arguably at too coarse a level to support the detailed decision making that both Heinrich [340] and Kjellen [444] argue must be informed by the monitoring of adverse events. Many regulatory organisations, therefore, publish more detailed information about the incidents and accidents that are reported by particular organisations. This helps to monitor the safety performance of those companies as well as their reporting behaviour. As we have seen, a noticeably low incident rate might indicate either a strong safety record or a poor reporting culture. For instance, Table 15.4 presents incident and accident statistics from Amtrack, the US National Railroad Passenger Corporation. Not only does this table provide an overall indication of incident frequencies, it also provides a more detailed breakdown of the causal factors associated with adverse events. As we have seen in Chapter 11 it can be difficult to ensure the consistency and reliability of such findings. For instance, it might be argued that changes in analytical procedures explain the marked rise in human factors related incidents rather than any underlying changes in operator intervention during adverse events and near miss incidents. It is important not to underestimate these analytical effects. For example, Cullen's analysis of the Ladbroke Grove rail crash concludes that a 'no blame' culture is an essential component of rail safety. However, he also acknowledges that this approach can encourage drivers to "accept blame in order to conclude the investigation as quickly as possible" [194]. This may help to explain why 85% of Signal passed at Danger (SPADs) are reported as driver error. Thic 'Cullen Paradox', therefore, implies that a 'no blaim' culture will make operators more likely to accept responsibility. These observations emphasise the importance of conducting further studies to validate results such as those shown in Table 15.4. By monitoring

Type	1997	1998	1999	2000	%age change 1997-2000
TOTAL ACCIDENTS & INCIDENTS	1,413	1,341	1,265	1,603	13.45
— Fatalities	117	120	105	131	11.97
— Nonfatal	1,328	1,180	1,161	1,412	6.33
TRAIN ACCIDENTS	84	89	85	148	76.19
— Fatalities	1	.	.	.	.
— Nonfatal	74	28	41	106	43.24
— Collisions	3	4	3	8	166.7
— Derailments	51	55	46	80	56.86
— Other	30	30	36	60	100.0
— Track causes	34	29	38	75	120.6
— Human factors	12	27	23	38	216.7
— Equipment causes	8	11	5	19	137.5
— Signal causes	.	.	1	1	.
— Misc. causes	30	22	18	15	-50.0
— Yard accidents	36	41	37	72	100.0
HIGHWAY-RAIL INCS.	176	170	181	202	14.77
— Fatalities	53	50	52	56	5.66
— Nonfatal	123	125	146	90	-26.8
OTHER INCIDENTS	1,153	1,082	999	1,253	8.67
— Fatalities	63	70	53	75	19.05
— Nonfatal	1,131	1,027	974	1,216	7.52
— Employee fatalities	3	2	0	0	-100
— Employee nonfatal	898	840	914	920	2.45
— Trespasser fatalities	57	67	51	70	22.81
— Trespasser nonfatal	32	30	25	18	-43.8

Table 15.4: FRA Accident/Incident Statistics, Amtrak, June 2001

changes in causal classification of incidents and accidents it is possible to gain important insights into the 'structural' weaknesses that can affect reporting systems.

Incident and accident frequencies can mislead the unwary in other ways. Previous statistics did not account for Amtrack's exposure to certain types of hazard. In particular, the data was not normalised for the relatively large number of rail operations performed by this company. In contrast, Table 15.5 provides normalised data for Amtrak and for the Grand Trunk Western Railroad. It is important that readers understand the ways in which incident frequencies are converted into normalised statistics. For instance, if we assume that normalised rail statistics are calculated by dividing the incident frequency by the number of train miles per year then a reduction in the incident rate might stem occur in several different ways. For example, it might be the result of a fall in the incident frequency with a stable number of train miles or of an increase in the train miles with a stable incident frequency etc. In practice, the FRA calculates the total accidents and incidents rate by multiplying the number of accident and incident reports by 1,000,000 and then dividing the result by the sum of train miles and hours. Similarly, the yard accident rate is the number of train accidents that occurred on yard track multiplied by 1,000,000 and then divided by the number of yard switching train miles. The 'other track' rate is the number of accidents that did not occur on yard track multiplied by 1,000,000 divided by the total train miles minus yard switching train miles. In contrast, the train accident rate is the number of train accidents multiplied by 1,000,000 divided by the total train miles. Highway-rail incident rate is the number of incidents multilied by 1,000,000 divided by the total number of train miles. The FRA's employee 'on duty' rate is the number of reported fatal and nonfatal cases multiplied by 200,000 and then divided by the number of employee hours worked. The trespasser rate is the number of reported fatal and nonfatal incidents, excluding

Type	1997	1998	1999	2000	%age change 1997-2000
Amtrak					
Total accidents/incidents	17.95	17.00	15.51	19.57	9.02
Train accidents	2.27	2.51	2.35	4.10	80.99
Yard accidents	18.39	19.70	17.78	34.60	88.19
Other track	1.37	1.44	1.41	2.24	63.48
Highway-rail incs.	4.75	4.80	5.01	5.60	17.90
Employee on duty	4.33	3.87	4.03	4.01	-7.20
Trespassers	2.40	2.74	2.10	2.44	1.57
Passengers on train	4.65	3.44	1.97	5.33	14.59
Grand Trunk Western Railroad Incorporated					
Total accidents/incidents	19.93	19.29	17.90	17.10	-14.2
Train accidents	4.42	3.91	4.05	3.71	-16.1
Yard accidents	15.79	6.32	9.99	5.95	-62.3
Other track	0.92	3.04	1.85	2.96	220.6
Highway-rail incs.	7.07	2.60	4.62	4.82	-31.8
Employee on duty	6.35	7.00	5.52	5.79	-8.80
Trespassers	0.53	.	0.39	0.93	74.82
Passengers on train	.	.	.	.	.

Table 15.5: FRA Normalised Statistics for Two Rail Operators

those associated with highway-rail incidents, multiplied by 1,000,000 and then divided by the total train miles.

The incident rates illustrated by Table 15.5 support several different monitoring activities. For instance, regulators can make detailed comparisons between the safety performance of companies with different operating characteristics. For example, Amtrak has a relatively stable incident rate for adverse occurrences involving trespassers. Grand Trunk Western has a relatively low trespasser rate which has increased rapidly in the period between 1997 and 2000. Such differences deserve further investigation. There may be operating changes that have increased Grand Trunk Western's exposure to these forms of incident. In which case, they may need to adopt the measures that Amtrak have taken to maintain their more stable rate. Alternatively, Grand Trunk Western's lower rate, even at the 2000 level, may suggest that Amtrak could learn more from their procedures. This example provides further illustration of the need to look beyond such statistics to understand the reasons for such differences.

Previous paragraphs have argued that it is important to consider both incident frequencies and the operating characteristics that are used to derive normalised statistics. Table 15.6, therefore, provides more detailed information about the employee hours, train miles and yard operations of Amtrak and the Grand Trunk Western Railroad. As can be seen, both companies reduced their total train miles between 1997 and 2000. The Grand Western's 14.2% reduction in accidents and incidents occurred when train miles only fell by 4.66%. In contrast, Amtrak's 9.02% increase in incidents and accidents occurred over a period when their train miles fell by 2.65%. As before, however, such analysis must be treated with care. Between 1997-2000, Amtrak increased their number of employee hours by 10.03% while those of the Grant Trunk Western Railroad fell by 9.60%.

Such caveats and complexities characterise the use of normalised incident frequencies as an indicator of the success or failure of incident reporting systems. It can be very difficult to associate particular trends with changes in the underling safety of an application. This problem is even more acute when metrics are used to identify the contribution that a reporting system can itself make to the operation of a safety-critical process. On the 30th November 1999, the UK Deputy Prime Minister, John Prescott, announced that the Confidential Incident Reporting and Analysis System (CIRAS) would be ext ended from the Scottish railway system to cover the entire network [100]. In

Type	1997	1998	1999	2000	%age change 1997-2000
Amtrak					
Train miles	37063760	35414704	36160704	36080704	-2.65
Yard switching miles	1,957,814	2,080,704	2,080,704	2,080,704	6.28
Employee hours	41663112	43480510	45399073	45840150	10.03
Passengers transported	20555107	21246203	21544160	22985354	11.82
Passenger miles	5.26888E9	5.32419E9	5.28868E9	5.57399E9	5.79
Grand Trunk Western Railroad Incorporated					
Train miles	5,657,394	5,376,050	5,190,349	5,393,620	-4.66
Yard switching miles	1,330,157	1,425,036	1,401,708	1,344,762	1.10
Employee hours	4,124,903	4,372,190	4,418,149	3,728,758	-9.60
Passengers transported	0	0	0	0	.
Passenger miles	0	0	0	0	.

Table 15.6: FRA Normalised Statistics for Two Rail Operators

the aftermath of the Ladbroke Grove crash he said that “I am pleased to say they have taken to heart everything that I asked of them in the wake of that terrible tragedy and today can announce concrete results on measures that can be taken now and commitment to a programme of action for longer-term projects... I repeat my pledge to the public that the industry will make rail travel even safer”. In spite of the perceived success of the CIRAS system, it is hard to demonstrate that Scottish railways have a significantly better safety record than other areas of the network. As we have seen, the region’s main operating company was one of ten that were warned by the Railways Inspectorate in June 2001 that they had not done enough to combat the problem of Signals Passed At Danger [112]. Such arguments suggest that there may well have been other motives behind the expansion of the CIRAS reporting system beyond the relative safety record of the company that operated it. For example, CIRAS’ original developers and operators [197] echo Clarke’s argument that ‘incident reporting might be viewed as an objective indicator of manager’s commitment to safety’ and that these ‘perceptions underlie a lack of mutual trust between staff and managers, which has implications for the fostering of open and honest communications within the network and for the development of a positive safety culture’ [170]. These sentiments are similar to those put forward by Cullen in his investigation into the Ladbroke Grove accident where he argues that confidential reporting systems would be unnecessary in an industry with a supporting safety culture [194]. Information about near-miss occurrences should be provided in an open manner without fear of subsequent persecution. Both arguments suggest that the potential utility of a reporting systems can be assessed in terms of the information that they provide about the safety culture in an industry.

In preparing this book, I have had many interviews with individuals who are involved in the development of the UK national rail reporting systems. In the course of these discussion, a number of criticisms have been raised about some of the arguments that are presented in the previous paragraph. These caveats illustrate the complex issues that arise during the monitoring of such applications. They also illustrate the way in which significant resources can be invested in the development of a reporting system even though there may be little consensus within an industry about the metrics that might be used to assess the success or failure of the system. Firstly, criticisms have been made about the statistics that were used by the HMRI SPAD report [112]. Secondly, performance in this area can be argued to have little connection with the information obtained from the CIRAS reporting system. Most ‘Signals Passed at Danger’ are observed by other rail personnel including signaling staff. They will, therefore, be notified by other means rather than the confidential incident reporting system. The success of the reporting system is, therefore, being assed in terms of safety-related incidents that it is not intended to address.

This section has identified a number of problems that frustrate the use of direct safety metrics as a means of monitoring the performance of incident reporting systems. These can be summarised as follows:

- there can be disagreement over the metrics that are used to assess the overall safety of complex applications. This creates problems when those metrics are, in turn, used to assess the contribution of a reporting system.
- it can be difficult to obtain data about the safety record of some applications even when there is agreement over the metrics to be used. Different jurisdictions can result in some data being withheld. Other organisations may under-report injuries and illnesses. This can make it difficult to assess the safety record of an industry which in tun complicates the use of direct metrics to assess the contribution of incident reporting systems.
- it can also be difficulty to identify normalising factors for statistical analysis. As we have seen, raw frequencies cannot easily be used to compare the performance of reporting systems in large and small organisations. There can be disagreements over the normalising factors to be used. It can also be difficult to collate the necessary operational statistics once those factors have been identified.
- incident reporting systems may only have an indirect effect on the metrics that are used assess the safety performance of an industry. This builds on Wright's arguments that the 'low severity' incidents described in reporting systems are very different in nature from the high-consequence accidents that are typically used to assess the overall safety performance of many industries [874].

These caveats have led some regulators to look beyond direct measures. Rather than assess the performance of a reporting system in terms of overall changes in the safety of an industry, more attention is paid to the indirect operational impact of each contribution. In other words, the success or failure of the system is assessed in terms of the different lessons that are learned from the incidents that are reported to it.

15.1.2 Indirect Feedback: Training and Operations

Previous sections have described how the Railway Group publishes regular summaries of rail safety across the UK rail network [691]. This reiterates the recommendations obtained from the CIRAS reporting system, mentioned above. For example, the survey published in October 2001 reminded managers 'at all levels in companies that are members of the Railway Group' that they should read the the publications from the national reporting system. In particular, they were advised to note the predominance of organisational problems in the incidents that were reported to the scheme. Most of these related to problems with rosters and shift patterns; 'short staffing is the most common perceived cause, and the most common consequence is fatigue'. The review also reiterated that poor communication by supervisors and management was a noted cause of many incidents. Rule violation was the most significant cause of what were described as 'workplace incidents'. The publication of this information is very significant. The majority of the Railway Group review is devoted to a statistical analysis of safety data, mainly focusing on the frequency of accidents and events that fall within the scope of a mandatory reporting system. The same approach is not used for the voluntary incident reporting system. Rather than providing statistics about contributions to the scheme or about the impact of recommendations on the frequency of accidents, the focus is on the lessons that have been learned from the system. This is an indirect approach because these lessons are intended to have a knock-on effect upon the other performance indicators.

The railway Group deliberately focuses on the high-level insights provided by the CIRAS incident reports. It does not identify particular recommendations and so there is a danger that they will have only a minimal effect on the recipients of the summary. Rather than directly measuring changes in accident and incident rates, it is possible to monitor the impact of a reporting system in terms of the changes that are made to operating practices. For example, UK reporting systems consistently revealed that track-side workers form the largest category of victims in rail related injuries and fatalities. In April 1995, these incidents led to the introduction of a relatively complex set of recommendations to segregate workers from trains. This involved a 'permit to work' scheme that ensured workers were either segregated from lines on which trains were running or that track workers were warned of approaching trains in time to move to a place of safety. Segregated worksites became known as 'green zones', while non-segregated worksites became known as 'red zones' [355]. These recommendations reduced but did not eliminate incidents involving track-side workers and so the HMRI started a further programme to review progress and to develop a strategy for future improvements. A questionnaire was developed to gather information about the effectiveness of previous recommendations based on the subsequent incident reports. These topics included the procedures used to monitor red and green zone working and the red zone risk assessment process. The results of these studies helped to identify further recommendations. In particular, it identified that some of the rail operators had provided misleading statistics when providing information about the normalising factors that, as we have seen, are important for the direct assessment of incident reporting systems:

> "A claimed 11% increase in green zone working, when analysed, represents a reduction in the proportion of green zone working because of a rise in the number of worksites (38% to 33% approx. over a twelve month period). The Railway Group Safety Performance Report 1998/99 has identified a need to provide information on an, 'exposed hours' basis for monitoring purposes" [355]

This illustrates the way in which regulatory and investigatory organisations can support investigations into the effectiveness of recommendations produced in response to previous incidents. These studies provide indirect insights into the utility of the reporting system. They can also yield additional recommendations that are intended to reduce the likelihood or mitigate the consequences of further incidents. Finally, they can also detect weaknesses in the way that a reporting system is currently being run. This is illustrated by the problems in reporting normalisation information, mentioned above. There are further examples of reporting systems being assessed in terms of the recommendations that they generate. For instance, the FRA's Switching Operations Fatality Analysis (SOFA) Working Group recently analysed 76 incident reports from January 1992 to July 1998 [241]. They also considered more limited FRA data from 1975 to 1991. The small total number of incidents and the varied circumstances of each event persuaded the Working Group that recommendations could not be based on formal statistical analysis. Instead, they used the incident data to devise 'five SOFA lifesavers'. These can be summarised as follows:

1. Notification to the locomotive engineer before fouling track or equipment. 'Any crew member intending to foul track or equipment must notify the locomotive engineer before such action can take place. The locomotive engineer must then apply locomotive or train brakes, have the reverser centered, and then confirm this action with the individual on the ground. Additionally, any crew member that intends to adjust knuckles/drawbars, or apply or remove EOT device, must insure that the cut of cars to be coupled into is separated by no less than 50 feet. Also, the person on the ground must physically inspect the cut of cars not attached to the locomotive to insure that they are completely stopped and, if necessary, a sufficient number of hand brakes must be applied to insure that the cut of cars will not move'.

2. Extra precautions when two or more train crews are working on the same track. 'When two or more train crews are simultaneously performing work in the same yard or industry tracks, extra precautions must be taken: C SAME TRACK. Two or more crews are prohibited from switching into the same track at the same time, without establishing direct communication with all crew members involved. C ADJACENT TRACK. Protection must be afforded when there is the possibility of movement on adjacent track(s). Each crew will arrange positive protection for (an) adjacent track(s) through positive communication with yardmaster and/or other crew members'.

3. Safety briefing. 'At the beginning of each tour of duty, all crew members will meet and discuss all safety matters and work to be accomplished. Additional briefings will be held any time work changes are made and when necessary to protect their safety during their performance of service '.

4. Proper communications. ‘When using radio communication, locomotive engineers must not begin any shove move without a specified distance from the person controlling the move. Strict compliance with ’distance to go’ communication must be maintained. When controlling train or engine movements, all crew members must communicate by hand signals or radio signals. A combination of hand and radio signals is prohibited. All crew members must confirm when the mode of communication changes’.

5. Paying proper attention to new crew members. ‘Crew members with less than one year of service must have special attention paid to safety awareness, service qualifications, on-the-job training, physical plant familiarity, and overall ability to perform service safely and efficiently. Programs such as peer review, mentoring, and supervisory observation must be utilised to insure employees are able to perform service in a safe manner’ [241].

These recommendations were published and then widely publicised within the US railway industry. There then followed a steady decline in switching incidents until in 2000 the FRA noted that the total number of switching-related deaths quickly exceeded those for 1999. These incidents raised questions about the working practices of crew members assigned to perform switching operations. They occurred on large and small railroads and included experienced employees with between two years to more than thirty years experience. This led to a review of the recommendations that had been derived from previous incidents. The FRA study concluded that most of the incidents ‘could probably have been prevented if all employees on each railroad had strictly followed the five recommendations of FRA’s Switching Operations Fatality Analysis (SOFA) Working Group and the applicable Federal and railroad company operating and safety rules to which they relate’ [241].

The previous paragraph illustrates the way in which the success of a reporting systems can be assessed in terms of whether particular recommendations might have prevented recent incidents. This approach has a strong appeal. As we have seen, there are few guarantees that regulators will accept the findings of reporting agencies. Similarly, companies may fail to implement the recommendations that are identified from previous incidents. The use of more direct reporting statistics ignores the impact that such factors can have upon the effectiveness of a reporting system. Indirect forms of analysis, similar to that presented by the FRA, serve to reiterate the lessons that might have been learned if these recommendations had been implemented.

It is also important to stress the limitations of these arguments. The FRA’s assertions about the effectiveness of the SOFA recommendations relies upon complex counterfactual arguments. More recent incidents would have been avoided had operating companies implemented the findings from previous incidents and accidents. Unfortunately, Chapters 10 and 11 have illustrated the dangers of this style of reasoning. In particular, it can be difficult to obtain evidence to support claims about the potential effect of recommendations that were not followed. The reiteration of well-known recommendations can also have a strong adverse effect if they are interpreted as needless reminders to ‘do better next time’ [409]. There is also a danger that by reiterating previous recommendations, regulators and investigators will fail to adequately consider the reasons why those findings were not followed in recent incidents. For instance, a study of incidents involving children near railways persuaded Administrator Jolene Molitoris that previous messages about the dangers of playing near railways had not been effectively communicated to the target audience. She, therefore, initiated the 1995 *Always Expect a Train* campaign using 270 television and cable markets, 673 radio markets and 194 publications [232]. As part of this work, a series of Public Service Announcements broadcast ‘deliberately graphic reenactments of motor vehicle-train collisions and railroad trespassing incidents, designed to grab the viewer’s attention’. A classroom teaching initiative as also created to embed safety-related information within multimedia resources on railway history and technology. These actions are instructive because they suggest a thorough re-evaluation of the way in which safety recommendations were communicated to the public. Such initiatives need not have been created if she had simply evaluated the reporting system in terms of whether previous recommendations might have prevented the incidents that were being reported. The recommendations publicised in these campaigns were essentially the same as those used in previous initiatives. In contrast, the successful implementation of the recommendations and of the incident monitoring system as a whole depended on the manner in which those recommendations were communicated to the target audiences.

These initiatives have been attributed with a 19% reduction in child-related rail 'casualties' [232]. Such statistics again raise the caveats and concerns that the previous section has raised about outcome statistics. However, the use of these figures to validate the revised recommendations is instructive because it illustrates the way in which most reporting systems are assessed both in terms of direct and indirect measures. The benefits of these systems are expressed both in terms of the particular insights that they provide and by the statistical reduction in severity or frequency rates. For instance, a fatal accident near Edson, Alberta, in early August 1996 forced Transport Canada to review their recommendations for avoiding runaway trains. In consequence, they encouraged a number of actions by the operating companies. These can be summarised as follows:

- training and education initiatives aimed at increasing employee and customer awareness of rules governing proper securement of cars;
- increased compliance monitoring by supervisors;
- increased inspection of derails to ensure proper application and positioning and to recommend locations where derails should be applied; and
- increased police monitoring of high vandalism areas [777]

The Alberta accident focussed the attention of the public and the rail industry on runaway train incidents. The high consequences of this incident led to demands for more direct evidence to demonstrate the effectiveness of these recommendations. It was insufficient simply to argue that the accident had led to the publication of the previous recommendations without also providing evidence that those recommendations were useful. Transport Canada, therefore, commissioned detailed comparisons between the number of runaway rolling stock incidents both before, between January and July, and after the publication of their recommendations, between August and December 1996. However, they anticipated that there would be a seasonal fall in the number of incidents in the winter months as the number of traffic movements fell. They, therefore, also compared this data with the number of incidents for corresponding months in 1994 and 1995. 42.3% of runaway rolling stock occurrences took place during the August-December period of 1996, compared with 44.4% during 1995 and 52.5% during 1994. These percentages are based on the total incident frequency for only the two periods that are considered in each year. The percentage of runaway rolling stock incidents that resulted in accidents in the August to December period fell from 60% in 1994 to 46.3% in 1996 and 46.5% during 1995. This created problems for the statistical analysis of the recommendations because 'the decrease in the percentage of uncontrolled movement incidents accounted for the entire decrease in runaway rolling stock occurrences that took place during the August-December period of 1996 when compared with 1994 and 1995 figures' [777].

15.1.3 Feed-forward: Risk Assessment and Systems Development

Kjellen distinguishes between four different levels of organisational learning [444]. These levels help to distinguish between different forms of metric that might be used to assess the performance of incident reporting systems:

1. *short-term learning in the workplace.* This describes immediate actions that are taken to address the direct causes of an adverse occurrence or near miss event. Kjellen argues that this form of learning only affects 'short-term memory'. In other words, any insights are likely to be forgotten as workers and supervisors move to new tasks or activities.

2. *long-term learning in the workplace.* This describes interventions that have a more sustained impact on operating practices within the particular work group or location where the incident took place. For example, they may result in the publication of revised operating guidelines or in documented modifications to particular pieces of equipment. Recommendations prevent recurrences but have limited scope and may not be effectively propagated throughout a factory or company.

3. *long-term learning in similar workplaces.* This can involve changes in the technical and administrative systems for the departments that are involved in an incident. Any recommendations will have a lasting effect and are likely to be propagated to similar departments in other areas of an organisation.

4. *long-term learning in management systems and norms.* These recommendations have profound effects on the way in which work is organised and managed. It can effect policy, goals and the specification of particular activities. The recommendations will affect most of the company and may have an impact on other organisations.

It can be argued that direct metrics, which focus on outcome measures, can be used to distinguish between these different forms of organisational learning. For example, level 1 recommendations may result in a short term fall in the accident and injury rates associated with a particular workplace. Level 4 changes will have a sustained effect on outcomes across many different sectors of an organisation. As we have seen, however, there are many factors that can confound the use of direct metrics to assess the impact of recommendations from a reporting system. In particular, the difficulty of obtaining reliable and appropriate statistical measures for safety improvements can be an obstacle to this approach. There are further problems. For instance, it can be difficult to distinguish between level 1 and 2 recommendations without careful monitoring of safety improvements over a prolonged period of time. Similarly, it can be difficult to distinguish between level 3 and 4 recommendations without reliable metrics for the performance of many different groups within an organisation. There are further problems. For example, it may take some time before particular recommendations have a discernible impact on outcome measures. These can be a delay before changes in the 'norms' and practices of senior management are effectively communicated into changes in operating procedures and acquisitions policy.

As we have seen, indirect measures do not focus on outcome metrics but, instead, concentrate on demonstrating the effective implementation of recommendations in the aftermath of an incident or accident. A range of further problems affect the use of indirect metrics as a means of assessing incident reporting systems. For example, commercial opportunities and regulatory intervention often force organisations to revise their working practices. These changes can lead to the introduction of new equipment and operating procedures. They can also invalidate many of the recommendations that were made in the aftermath of previous failures. In such circumstance, it can be difficult to distinguish between situations in which those recommendations have been 'forgotten' and situations in which previous recommendations no longer apply to present working practices. These problems are compounded by the difficulty of indirectly monitoring long-term changes in management practices. It is easier to identify level 1 changes than it is to demonstrate the effective implementation of level 4 recommendations. New piece of equipment and revise manuals are more tangible than changes in management 'norms'. Attitudinal questionnaires often focus on short-term effects and are subject to a host of biases [342]. This makes it difficult to interpret the results of such surveys, especially in the aftermath of safety-related incidents.

A number of authors, including Benner [73], have argued that evidence of long-term 'organisational' learning can be obtained by examining the impact that adverse events have upon risk assessment practices. This approach addresses many of the criticisms of direct and indirect metrics that were introduced in the previous paragraphs. For example, risk assessment practices provide a useful measure of management norms because they have a direct impact upon the allocation of finite resources. Risk assessments reflect the priorities associated with development and maintenance activities and hence indicate operational concerns at higher levels within an organisation [188]. The outcome of risk assessment procedures can also be used to predictive potential safety problems. In other words, the priorities derived from risk assessments helps to identify those areas that managers believe will pose the greatest threat to the future safety of an application. This offers the opportunity for analysis to determine whether incident statistics actually support those priorities. In contrast, direct metrics provide information about the post hoc success or failure of previous operational decisions.

A number of practical problems complicate the use of risk assessment metrics to assess the performance of incident reporting systems, For instance, many local incident reporting systems are isolated from the revenue streams that are necessary to fund large-scale safety improvements [417]. In contrast, they must fund the implementation of safety recommendations from the savings that are made through previous recommendations. This approach is intended to ensure that incident reporting systems are well integrated with wider 'lessons learned' applications. It also ensures that the reporting system is self-funding. Unfortunately, such practices also isolate the reporting system from normal risk assessment practices within the rest of the organisation. Important insights about the causes of previous incidents and accidents may not be communicated to those individuals who have the greatest influence on future acquisitions. In contrast, many other reporting systems are explicitly integrated into risk assessment systems. For instance, Transport Canada's guidance on the development of rail safety management systems identifies three stages to the risk management 'process' [780]:

1. Identification of Safety Issues and Concerns
   The first stage of risk management explicitly focuses on gathering 'input from incident/accident investigations and safety data collection and analysis'.

2. Risk Estimation
   The information from the first stage of the process is then analysed to assess the probability and severity of a potential hazard using either qualitative or quantitative techniques. Quantitative estimates can 'sometimes be developed from safety performance data, illness and injury records'. Transport Canada do, however, note that probability estimates based on 'historical data assume that future conditions will mirror those of the past'. If there is no relevant incident data then more qualitative techniques, such as event-tree analysis should be used to generate risk estimates. Event tree analysis enumerates the outcomes from a given event to map out the likely sequences of consequent events. For each event, analysts can consider the consequences of safety systems failing or succeeding in their specified function. Probabilities can then be associated with each path through what can be thought of as a form of decision tree [838].

3. Risk Evaluation
   The final stage of the management process determines which risks are tolerable, tolerable with mitigation and or unacceptable. These decisions should be guided by classification methodologies based around the Risk Assessment Matrices described in Chapter 12.

The Railway Safety group exploits a similar management approach to risk across the UK rail network [691]. Their three stage model manages risk by 'understanding the relationship between precursors and incidents... then by measuring the precursors and finally by applying action to the areas identified. This risk management process is supported by a Precursor Indicator Model (PIM). This relies upon 16 measures that were identified through the analysis of previous accidents and incidents. Table 15.7 enumerates these precursors. It also provides a percentage indicator that is intended to represent the 'risk' associated with each contributory factor to catastrophic rail accidents. Arguably a better description would be the percentage of major accidents in which analysts identified these precursors. The Railway Safety group interprets table 15.7 as providing a 'severity weighting' for precursors by arguing that 'a third of all injuries from train accidents are caused by category A SPADs'.

The historic severity weightings identified in Table 15.7 can be used to address some of the limitations of direct metrics. In particular, it provides a means of assessing the safety of complex systems that suffers very few catastrophic failures. We begin by pairing the severity assessments of each precursor from Figure 15.7 with the number of times that the precursor has occurred in the time period under consideration. There pairs can be used to construct a vector of the form $(\mathit{frequency}_1, \mathit{historic\_weighting}_1;\ ...;\ \mathit{frequency}_n, \mathit{historic\_weighting}_n)$. It is important to note that precursor frequencies can be obtained even though these incidents may not have led to a major accident. However, the Railway Safety group can use this information to calculate the overall risk of a major accident in the following way:

$$\mathit{system\_risk\_assessment} =$$

Precursor	Proportion of Train Accident Risk
Category A SPADs	32.84%
Level crossing misuse	22.84%
Track quality	12.86%
Irregular working	8.31%
Rolling stock failures	8.00%
Environmental factors	6.07%
Vandalism	2.91%
Structural failures	1.54%
Train speeding	0.98%
Level crossing failures	0.90%
Irregular loading of freight trains	0.83%
Wrong-side signaling failures	0.36%
Non-rail vehicles on line	0.27%
Possession irregularities	0.15%
Hot axle box	0.13%
Animals on the line	0.04%

Table 15.7: UK Railway Safety Group's Precursor Indicator Model (PIM)

$$\sum_{i=1}^{n} \mathit{frequency}_i.\mathit{historic\_weighting}_i \tag{15.1}$$

As mentioned, the procedures used by UK Railway Safety show how incident and accident data can be used to replace direct measures of system safety for infrequent, high-consequence events. The model exploits failure information in two ways. Firstly, the weightings associated with different precursors depends on the frequency of their observation in previous accidents and incidents. This ensures that the relative importance of those weightings will change as different failures become more or less significant to the overall safety of the rail system. It is, however, also important to ensure that these weightings are derived from a relatively large sample of previous failures to ensure that adequate attention is paid to long term problems as well as more immediate changes in the precursors to incidents and accidents. Secondly, the frequency of particular precursors is introduced into the calculation of overall system risk. These precursors are directly identified from recent incident and accident reports so that any calculations reflect more immediate changes in the performance of underlying systems. The UK calculations for October 2001 reflect reductions in the frequency of track quality faults, wrong-side signaling failures and train speeding. They also reflect increases in level crossing misuse, irregular working, vandalism, level crossing incidents and rolling stock failures. These changes in frequency combined to create a slight increase in the overall accident risk in the first half of 2001/2002 [691].

Although the PIM approach illustrates both a comprehensive and effective integration of risk assessment and incident reporting, it is possible to identify a number of potential problems. As we have seen, the use of previous accident information to identify incident precursors will only provide reliable risk assessments if future incidents are similar to those that have occurred in the past. It is, therefore, essential that the components of Table 15.7 be reviewed at regular interval to ensure that they do not exclude potential causes of future failure. A further problem is that the PIM approach fails to distinguish between the different levels of severity that are associated with catastrophic accidents. This is significant because many safety objectives are expressed in terms of the number of fatalities per train miles. The UK objective is 0.3 fatalities per million train miles by 2009; in 2000-2001 the annual moving average was 0.59 while in 2001-2002 it was given as 0.52 [691]. The Rail Safety group recognise this problem and, therefore, calculate a weighting for outcomes based on previous accidents. These are expressed in terms of 'equivalent fatalities per 10 train accidents' over a specified period of time. t is important to emphasise that this metric relates to different types of accidents and not the precursors, mentioned above, that lead to those accidents. Current weightings based on a 16-year interval are given in Table 15.8. The relatively low weighting associated with 'buffer stop collisions' has led some analysts to question whether they should be included in the significant train accident statistics.

Type of Accident	Consequence Weighting
Passenger collisions	7.331
Non-passenger collisions	0.777
Passenger derailments	0.927
Non-passenger derailments	0.005
Buffer-stop collisions	0.162

Table 15.8: UK Railway Safety Group's Significant Train Accident Weightings

Table 15.8 illustrates the way in which safety improvements can have a mitigating effect on the outcomes of adverse events. It does not, however, illustrate the procedures by which precursors are associated with particular types of outcome. This is far from straightforward. Recall from Chapter 11 that Bayes theorem considers the probability of a given hypotheses, B, in relation to a number of alternative hypotheses, B_i where B and B_i are mutually exclusive and exhaustive:

$$Pr(B \mid A \wedge C) = \frac{Pr(A \mid B \wedge C).Pr(B \mid C)}{Pr(A \mid B \wedge C).Pr(B \mid C) + \sum_i P(A \mid B_i \wedge C).P(B_i \mid C)} \tag{15.2}$$

Bayes' theorem can be used to assess the probability of a particular factor or precursor, B, causing a failure given that an incident report has identified that causal factor, A. Suppose we examine reports of previous buffer stop collisions to determine whether or not they were caused by train speeding, B. It is unlikely that we will have complete confidence in the intuitive causal analysis of every investigators. We, therefore, conduct a quality control exercise that performs a more detailed causal analysis for a sample of recent reports. This indicates that there is a false positive rate of 4%. In other words, 4% of the reports argue that the driver was speeding when they were not. We might also conclude that the false negative rate is 3%. This is the percentage of reports that suggest speeding was not the cause when subsequent investigations revealed that it was. The reports were, therefore, 96% accurate for incidents in which the buffer stop incidents were caused by train speeding. They were, in contrast, 97% accurate for incidents that were not caused by this precursor. To simplify the exposition, we assume that further analysis reveals that 1% of incidents were caused by train speeding. Table 15.7 provides the more accurate figure of 0.98%. We can use the previous formalisations to determine how likely it is that train speeding caused a buffer-stop incident given that a report identifies this as a cause of an incident. The probability of train speeding having caused the buffer stop is less than 20% based on a positive incident report! The following formulae adopt the convention of including the context C for the reasons given in Chapter 11. Dembski uses a variant on this example to demonstrate the ways in which people can fail as 'intuitive probabilists' (see [200], pp 83-84). These apparent failures are so deeply engrained that I remain unconvinced by aspects of his argument even if I can agree with the underlying mathematics!

$$\begin{aligned} & Pr(B \mid A \wedge C) \\ & = \frac{Pr(A \mid B \wedge C).Pr(B \mid C)}{Pr(A \mid B \wedge C).Pr(B \mid C) + Pr(A \mid \neg\, B \wedge C).P(\neg\, B \mid C)} && (15.3) \\ & = \frac{(0.97).(0.01)}{(0.97).(0.01) + (0.04).(0.99)} && (15.4) \\ & = 0.1968 && (15.5) \end{aligned}$$

A number of further practical issues complicate the use of incident data within the PIM approach to risk management. In particular, it is unclear how specific interventions by particular companies are to be identified from high level, industry-wide statistics. One possible mechanism is through the enforcement actions that are recorded in the Safety Group reviews. These record the date, location and nature of each violation that triggered a prosecution. Examples include 'all reasonably practicable measures have not been taken to reduce the risk of signal XXXX being passed at signal danger' and 'redundant coach left on disused track beneath bridge being vandalised and used by children'. These incidents are fed into the calculation of precursor frequencies mentioned in previous sections. The focus is, however, on the correction of violations rather than on the pro-active interventions that might further reduce potential risks.

A detailed analysis of the impact of particular mitigation or risk reduction measures arguably lies beyond the scope of the UK Safety Group's periodic reviews. Specific guidance is included within 'Focus Area' publications on reducing SPADs, Trackworker Safety, Trespass and Vandalism. Transport Canada provides a further example of the use of incident reporting data to inform risk assessments. As mentioned, their three stage risk management process is similar to that advocated by the UK Railway Safety group. This approach was employed to support the analysis of six incidents in which hot bearings had led to 'burnoffs' on the Canadian rail system. Hot journal bearings occur when inadequate wheel bearing lubrication or mechanical flaws cause an increase in bearing friction. If undetected, the resulting rise in bearing temperature can lead to a bearing burnoff which can cause a derailment. The analysis began by assessing the frequency of previous incidents. An average of approximately 10 derailments per year were linked to burnoffs between 1987 and 1993. In other words, there were only 2 burnoffs per billion freight car miles. The consequences of these events were also assessed. These incidents did not result in any fatalities; 'no passenger or member of the public has been fatally injured because of any main track derailment for over 10 years' [776]. The investigation assumed an average cost of $250,000 per derailment in terms of damage to rolling stock and infrastructure. This analysis of the risk associated with these incidents was then used to determine whether or not to invest in a number of detection and mitigation techniques. For example, the Canadian rail system already had a number of 'hot box' detectors. One proposal suggested that the number of these detectors be doubled so that the distance between consecutive units would be reduced to around 12 miles. This would involve 400 new single track and 200 double track installations at an initial cost of $90 million with a further $5 million per year for maintenance. Six detailed investigations led investigators to conclude that the additional detectors would, at best, have prevented half of these derailments. This would have brought the annual average number of burnoffs down from to 5 per year saving around $1.25M per year given the cost estimates for each previous derailment The investigators concluded that 'spending $90 million plus $5 million per year on additional detectors to save some 5 burnoffs per year without any evidence that lives would be saved or injuries prevented is clearly not a beneficial use of society's resources, is out of line with risk management expenditures in other areas, and is not a recommended course of action for the Railway Safety Directorate or the industry to pursue' [776].

The previous example illustrates a number of links between such risk assessments and the monitoring of incident reporting systems:

1. *issue identification.* Firstly, incident reporting schemes help to trigger risk management activities. Evidence from previous derailments justifies the initial investigation of bearing burnoffs as a causal factor. In this way, the incident reporting system 'earns its keep' as a necessary part of a safety management system even though it may not be possible to identify suitable metrics to support the performance of the scheme in issue identification.

2. *data sufficiency.* Secondly, the previous example illustrates the way in which risk assessment activities provide a means of assessing whether incident investigations yield sufficient insights into the causes of adverse events. In particular, the Transport Canada investigators make use of the counter factual argument that doubling the number of hot box detectors would only have prevented about half of the annual number of derailments from bearing burnoffs. Such an analysis depends upon sufficient information being available to support their conclusions. If such evidence had not been available then the risk assessment would have been seriously

flawed and the investigatory process would have been modified.

3. *investment savings.* One of the side-effects of integrating incident reporting into risk management procedures is that it provides monetary assessments of the strategic value of information that is provided about previous failures. In the previous example, it can be argued that this data helped to avoid investing more than $90 million in a scheme that would have yielded limited safety benefits. Sadly, such superficially appealing arguments can be challenged in a number of ways. They assume that the costs associated with future incidents will continue to similar to those incurred by previous failures. The $90 million costs might, however, appear to be justified if a future incident resulted in multiple fatalities.

The integration of incident reporting systems into risk management procedures provides a powerful justification for the investment that is needed to elicit and analyse information about previous failures. The previous list, therefore, shows how the utility of a reporting system can be indirectly assessed in terms of the support that it provides for risk assessment. Such approaches are unlikely to be sufficient. In particular, they provide relatively little information about the effectiveness of the reporting system in eliciting information about adverse events. Similarly, these techniques cannot easily be used to address the problems of intra and inter analyst reliability that have been identified as a potential limitation for direct, indirect and feed-forward metrics. The following section, therefore, focuses on the process measures that can be used to assess the performance of the reporting system itself rather than the utility of the information that it produces.

15.2 Process Measures

It is important to monitor the costs that a reporting system incurs as well as the benefits that it delivers in terms of safety improvements. For example, the Aviation Safety Reporting System spends about $3 million annually to analyse roughly 30,000 reports, at about $100 per case. These techniques would cost almost £50 million annually if the same techniques were applied to the 850,000 adverse events in the UK National Health Service [480]. Such figures illustrate the importance of retaining a close control of the management of incident reporting systems. It is difficult to obtain similar estimates for national rail reporting systems. The UK Health and Safety Executive argued that 'trials carried out in Scotland since 1996 (involving ScotRail, GNER, and Virgin (N)) indicate that the CIRAS confidential incident reporting system improves incident reporting as well as being financially beneficial to the companies concerned' [325]. They did not published detailed figures to support this argument and several operating companies expressed concerns about the financial overheads associated with voluntary incident reporting. Their concerns echo criticisms voiced about the development of national reporting in Australian railways. The Booz, Allen and Hamilton report recognised that individual rail operators, infrastructure managers and regulators gather data to monitor their own performance over time [55]. It was the lack of 'consolidation, consistency and analysis' at a national level that gives the greatest cause for concern because 'the industry is not yet convinced that these processes are consistently applied or completely appropriate'. In consequence, the proponents of national reporting systems have been forced to monitor the operation and management of their schemes and not simply the safety-related information that they produce.

15.2.1 Submission Rates and Reporting Costs

A number of crude measures can be used to assess the cost effectiveness of a reporting system. For example, the total investment in any scheme might be divided by the number of reports that are received each year. There are a number of potential benefits from using submission metrics to support the monitoring of incident reporting systems. In particular, this can help to identify the problems of under-reporting that were studied in Chapter 5. As the FRA note, employees may even neglect medical treatment rather than expose themselves to workplace harassment:

> "FRA has become increasingly aware that many railroad employees fail to disclose their injuries to the railroad or fail to accept reportable treatment from a physician because they wish to avoid potential harassment from management or possible discipline that is sometimes associated with the reporting of such injuries. FRA is also aware that in some instances supervisory personnel and mid-level managers are urged to engage in practices which may undermine or circumvent the reporting of injuries and illnesses." [233]

There are a number of problems with using submission rates as a metric for overall system performance. Most reports are received from a relatively small section of the workforce in many industries. Increases in the numbers of contributions from these employees can mask significant under-reporting in other areas. This point is illustrated by Table 15.9, which presents statistics on fatal and non-fatal injuries on US railways [243]. The majority of reports are filed by workers 'on duty' and under the employment of a rail operating company. In contract, there are relatively few reports form rail contractors even though they make up a significant proportion of the workforce employed in maintenance and infrastructure projects in this industry. 119 'workers on duty' were killed on US railways between 1997 and 2000. During this period, they suffered 33,738 non-fatal injuries. This yields a Heinrich ratio of 283.51 non-fatal injury reports per fatality. In contrast, 31 contractors were killed between 1997 and 2000. During this period only 1466 injuries were reported at a Heinrich ratio of 77.29. Contract workers might be less likely to be involved in incidents than other workers and hence we might expect a lower ratio of non-fatal injuries to fatalities. There is, however, considerable evidence to the contrary [342]. Contract workers are often less well trained and briefed on their operating tasks than full-time employees. They are also less easily integrated into working groups when they may be moved between operational responsibilities more frequently. In consequence, it can be argued that the difference in ratios illustrates a problem of under-reporting of non-fatal injuries amongst this section of the workforce. Simply dividing the overall number of non-fatal incident reports by the annual expenditure on a reporting scheme would fail to identify such structural problems. Similarly, Table 15.9 illustrates the difficulty of identifying an underlying pattern in the submission data for non-fatal incidents. Short-term reductions, for example in passenger-related incidents between 1997 and 1999, can be offset by increases elsewhere, for instance in employee on duty submissions.

	Fatalities				Nonfatal Conditions			
	1997	1998	1999	2000	1997	1998	1999	2000
Worker on duty (railroad employee)	37	27	31	24	8,295	8,398	8,622	8,423
Employee not on duty	.	2	.	1	263	219	216	286
Passenger on train	6	4	14	4	601	535	481	658
Nontrespasser	362	324	302	332	1,517	1,201	1,307	1,264
Trespasser	646	644	572	570	728	677	650	606
Worker on duty(contractor)	6	2	2	.	213	237	172	183
Contractor(other)	5	3	10	3	121	143	212	185
Worker on duty(volunteer)	.	.	.	.	3	11	4	6
Volunteer(other)	.	.	.	.	3	3	1	2
Non-trespasser, off rr prop	1	2	1	3	23	35	35	30

Table 15.9: FDR Rail Incident Reports by Worker (1997-2000)

One way of avoiding such problems is to attempt to increase contributions from particular sections of a workforce while maintaining or reducing the overall cost of operating a reporting system. This raises further problems. As we have seen in Chapter 2 short-term changes in submission rates can occur independently of changes in the safety of an underlying application. Awareness arising campaigns can elicit large numbers of 'low risk' contributions from a minority of the target workforce without encouraging the mass of their colleagues to report on the more serious incidents that are masked from a reporting system. Such campaigns are often costly and the effects that they achieve can be very short-lived [444, 342]. There are a number of further ways in which submission rates can fail to provide accurate indicators of the underlying safety of an application process. Such metrics are profoundly affected by changes in the criteria that are used to identify particular types of incident. For instance, the UK Railway Group widened the definition of SPAD severity Category 3 in May 2000. The new definition increased the number of incidents falling into this category. It provided a larger data group and was, therefore, argued to increase the opportunity for more meaningful analysis. Such changes created the need to revise previous data for category 3 SPADs in order to reflect the new definition and provide a consistent basis for comparison. This reclassification also illustrates how structural changes on a reporting system can have knock-on effects both on efficiency metrics and more direct forms of risk assessment. The revision indirectly increased the efficiency of the reporting system in terms of the severity of incidents being analysed within a particular budget. It also reduced overall system safety measured in terms of risk metrics, including the Precursor Indicator Model mentioned in the previous section.

The complexity of using submission and frequency metrics to assess the performance of reporting systems has persuaded many managers to concentrate on monitoring the costs of their schemes. This approach is justified by a series of interviews and focus groups that the FRA conducted to identify concerns about the regulation of the US rail system [245]. This study was intended to identify the influence of corporate culture on compliance with railroad operating rules. Part of this work focused on the attitude of operating companies to the opportunities provided by incident and accident reporting. There was a degree of skepticism about whether the costs incurred in analysing 'near accidents' could be justified by the potential insights that they provided. The rare nature of these events implies that 'one must make any number of assumptions' to identify the potential root causes and that this 'inevitably reduces the level of certainty' Many in the industry were concerned about the large numbers of near-miss events that must be investigated to gain limited insights about a small number of actual incidents; 'analysing near-incident data substantially increases the population data set from which to study'. In particular, the FRA study found that the 'systematic analysis of probable cause' for near miss incidents involving the safety conduct of locomotive engineers 'is seldom conducted' under Federal Regulation 49 CFR 240.309 [245]. To address these costs issues, the FRA report proposed that greater use be made of automated data analysis tools from reporting schemes in other industries. These tools have been reviewed in Chapter 14. In particular, they argued that Internet and Intranet technologies should be exploited to reduce the costs associated with the analysis of near-miss incidents.

Some of the concerns identified by the FRA stem from the difficulty of assessing the costs associated with running a reporting system. It might seem relatively straightforward to account for the fixed costs that are associated with infrastructure items, such as computer hardware. Most of this equipment serves several different purposes. Incident reporting software is often integrated into other aspects of a safety management system. It can, therefore, be difficult to distinguish the costs of running such a system from the wider overheads associated with risk analysis and assessment [780]. Similar comments can be made about the difficult of auditing variable costs. Many reporting systems rely upon the involvement of volunteers who combine incident analysis with more a more direct operational role [658]. Cost estimation for incident reporting is further complicated by the consequential overheads that are associated with some investigations. For example, the UK Health and Safety Executive make a charge for inspectors' time under The Health and Safety (Fees) Regulations 2001. Employers will be charged for the 'investigation of activities or workplaces where HSE becomes aware of an incident which has caused or is liable to cause injury to persons; and related enforcement work... including the preparation and serving of improvement or prohibition notices; assessing and issuing exemptions" [329]. Employers' scope for cost reduction is further constrained by the need to support a reporting system as a condition of operation. For example, the Safety Case for operating the London Underground contains a specific commitment to operate such a system; 'employees can raise health and safety concerns either formally with management or informally via their employee representatives... London Underground Ltd intend to join the CIRAS confidential reporting system when it is eventually rolled out across the national railway network'

[333].

It can be just as hard for regulators to justify and account for their expenditure on incident reporting systems. For instance, the FRA required $106,855,000 for 'safety and operations' in 2001. In 2002, they requested $120,583,000 while in 2003 this had risen to $122,889,000. It can be difficult to break these figures down to identify the individual sums spent on the diverse range of incident and accident reporting systems that are supported by the FRA within the US Department of Transportation. However, the 2003 budgetary request provides an important insight into the strategic importance of these systems when it justifies the additional expenditure. This additional money is intended to fund 20 new safety field inspectors because 'the number of railroad issues facing FRA is increasing and becoming more complex' [242]. It is hoped that these posts will support the 'elimination of transportation-related deaths, injuries, and incidents'. The FRA's request does not monitor the efficiency of their reporting system in terms of the cost per submission, as suggested in previous paragraphs. They do, however, cite a number of additional statistics to support their request for additional funds. These are similar to the normalising factors that are used in the calculation of direct measures for reporting system performance. The budget request focuses on the 55% increase in freight traffic since US deregulation in 1980 [242]. They also stress the unfortunate rise in rail passenger fatalities and injuries from approximately 500 per year over the past decade to over 660 in 2000. The increase in rail traffic and in adverse event is, therefore, used to justify increased regulatory expenditure on field investigations.

To summarise, performance metrics focus more on the operation of the reporting system than the direct measurement of 'system safety'. For example, the performance of a reporting system be assessed in terms of the number of incidents that are analysed within a specified budget. The efficiency of the system can be improved either by budget reductions or increases in the number of incidents that are handled by the system. Unfortunately, a number of pragmatic and theoretical concerns affect the use of such monitoring techniques. Raw data about the number of submissions made to a reporting system can be very misleading. For instance, they can hide under-reporting by particular groups. Any increase in reporting frequency might, therefore, yield few marginal benefits if those reports stem from communities that are already well represented with the system. It is also possible to distort submission statistics by changing the definition of what does and does not fall within the scope of the system. Alternatively, efficiency can be assessed in terms of savings that can be made from the operation of a reporting system. Opportunities for cost reduction may, however, be constrained by regulatory agreements that require the operation of particular schemes. The increasing complexity of many application processes also makes it difficult to reduce the costs associated with incident reporting. Both the FRA [239] and HMRI [319] have been forced to find and fund an increasingly diverse range of skills during the investigation of many recent failures. Staff costs represent the greatest single investment in most reporting systems. In consequence, managers have begun to refocus their monitoring activities on the performance of their investigators rather than on more direct metrics that can be both difficult to gather and harder to interpret.

15.2.2 Investigator Performance

Many regulatory organisations have developed detailed guidance for the investigation of adverse events and near-misses. For example, Transport Canada include such advice in their recommendations for the development of railway safety management systems [779]. Operating companies must develop 'procedures for internal and external accident and incident notification and reporting, including third-party reporting; procedures, formats and approaches (e.g., site protocol) for investigations (e.g., environmental, employee injuries, transportation of dangerous goods); a formal link to the risk management process; and procedures for reporting and documenting findings, conclusions and recommendations, and for ensuring implementation of recommendations and corrective actions'. The Safety Management System guidelines go on to argue that most train accidents can be prevented and that investigators must, therefore, identify ways of providing both 'immediate protection' and 'long-term correction'. Examples of an immediate action include the introduction of a 10 miles per hour temporary speed restriction at the site of a track geometry defect or a 40 miles per hour speed restriction on a type of car that appears to be unstable at higher speeds. These immediate protective actions must be implemented by the Investigating Team before operations are resumed. Long term 'corrections' reduce the likelihood of a similar train accident recurring in the future. Examples include the 'accelerated removal of straight plate wheels and the overhaul of trucks on a specific class of car' [779].

Such guidance reflects the way in which many regulatory and investigatory agencies have sought to support the work of incident investigators. In particular, it is typical of the way in which general recommendations are not supported by more detailed assessment criteria that might be used to monitor the performance of particular teams and individuals. This is a significant concern which is shared by many of the organisations that operate incident reporting systems [423]. Often investigators are domain experts, drawn from diverse operational areas of the industry that they are helping to protect. This offers numerous benefits in terms of their detailed understanding of industry practices. It also creates significant weaknesses because many investigators have only a rudimentary understanding of the more detailed causal analysis techniques pioneered by NASA [571], the US Department of Energy [207] and the NTSB [87]. There are further problems. The previous background of an investigator can have a powerful influence on their likely findings. Lekberg describes the correlation between an investigators area of expertise and the likely results of a causal analysis [484]. The problems of frequency and recency bias have also been described in Chapter 11. In particular, the fact that investigators may have already spent many years within an industry can imply a lack of understanding of more recent technical innovations. For example, new insights in the field of human factors are often slow to inform the conduct of incident investigations [699].

There are relatively few published studies into the performance of investigators. This is a significant barrier to the future development of incident reporting systems. One consequence is that many of the organisations that operate these schemes are concerned about potential inadequacies in their analysis and interpretation of individual reports. Unfortunately, the lack of previous published work in this area has created a general reluctance to assess or otherwise measure the extent of the problem. There are several exceptions. For example, the UK HMRI conducted a recent analysis of the work that investigators performed in analysing the causes of SPAD incidents. The HMRI enquiry 'looked at the results of several SPAD investigations carried out in accordance with procedure GO/RT3252 in each Railtrack zone, and it appeared that in some cases greater emphasis was placed on completing a multi-page form than getting to the root cause of the SPAD incident' [349]. They describe examples in which the same signal had been passed at danger on repeated occasions and yet the cause had not been established. It could be argued that this reflects a lack of evidence available to any investigation, however, the investigatory procedures stressed the need to reach some form of closure for each enquiry. In other cases, investigators had failed to follow through their analysis of an adverse event. For example, one investigation concluded that the driver had an 'unsuitable temperament' for driving suburban trains. He was, therefore, barred from driving these services. The HMRI inspectors argued that the investigation should have gone on to deal with the root cause of the incident which they interpreted to be the 'inadequacy of the measures for assessing the competence of drivers' [349]. The inspectors also found that SPAD investigators had difficulty in distinguishing between some of the causal categories identified in the supporting documentation. In particular, it was often unclear whether an incident was caused by 'misjudgement' or 'disregard' as it implied an assessment of driver intentions. The HMRI inspectors cite the example of a SPAD that occurred in poor weather conditions where the driver made every effort to stop the train before the signal. They argued that this was incorrectly classified as 'disregard' rather than 'misjudgement'.

Arguably the most straightforward means of assessing the performance of incident investigators is to introduce self-monitoring throughout team-based enquiries. This approach has been widely adopted and is often modelled on the 'Go Team' procedures that were initial developed by the NTSB during the 1960s and 1970s. Each Go Team is led by an Investigator-in-Charge who is a senior investigator with several years of NTSB and industry experience. Each member of the team is responsible for a defined aspect of of the investigation. For example, the 'operations specialist' will reconstruct the history of the incident including the crew members' duties for the period before the incident. The 'structures expert' will document the accident scene. They will analyse any wreckage and will also calculate impact angles and speeds. The 'human performance specialist' will make a study of crew performance and all before-the-accident factors that might be involved in human error, including fatigue, medication, alcohol. They will also consider the possible role of drugs, medical histories, training, workload, equipment design and work environment [617]. Locomotive engineers, signal system specialists and track engineers head working groups at railroad accidents. Each of these specialists heads a working group in their area. The members of the working groups are drawn from 'interested parties'. The party scheme is the key to the NTSB's efficiency; it investigate approximately 2,500 incidents per year with only 400 employees. The NTSB designates other organisations or corporations as parties to the investigation. These parties must provide specific expertise to the investigation. There is considerable freedom about who might provide such assistance; the only exclusion is that 'persons in legal or litigation positions are not allowed to be assigned to the investigation' [617].

Not only does the party system increase the efficiency of the Board's full-time staff. It also provides an important means of cross validation and of monitoring the quality of the investigation process. The head of each working group prepares a factual report and each of the parties in the group are asked to verify the accuracy of the report. Unfortunately, the formal procedures and mechanisms that support NTSB investigations are resource intensive. Most incident reporting systems lack the resources necessary to finance the involvement of more than one or two investigators in the analysis of adverse events. There are further problems. For example, the NTSB are an independent organisation that operates across several different industries. It is difficult to see how such a multi-party architecture might be used by a proprietary reporting system which focuses on adverse events within a single commercial organisation.

Fortunately, a range of alternative techniques can be used to monitor and support the performance of incident investigators. For instance, the ATSB was formed with the specific aim of pooling expertise in transport safety. It created a unified framework for the Bureau of Air Safety Investigation, elements of the Federal Office of Road Safety and the Marine Incident Investigation Unit and a new Rail Safety Unit. The intention was to 'make safety investigations even better as a result of sharing resources, ideas and techniques' [54]. The Bureau encourages investigators to move between incidents in different modes of transport. This ensures that key skills acquired in the investigation of rail incidents support the analysis of aviation or maritime incidents and vice versa. For instance, the ATSB reports that air safety investigation techniques were applied to the freight train collision at Ararat in November 1999 [54]. This exchange of expertise can also be seen to support the monitoring of individual investigators who must move beyond the immediate area of their core expertise. The Canadian Transportation Safety Board have similar objectives. Not only have they sought to improve the exchange of expertise between different investigation modes, they have also attempted to increase consistency in the resources that are available to investigators. In particular, they have developed a multi-modal Statement of Requirements that emphasises the importance of recorded information for investigative purposes in aviation, marine, rail, and pipeline incidents. The intention is to provide investigators with sufficient information to ensure that their reports achieve a level of 'reliability, comprehensiveness and timeliness' regardless of the mode of transport [88]. These multi-modal approaches are innovative and challenging. It remains to be seen whether they will realise the benefits that their proponents anticipate. They do not, however, provide a panacea for incident reporting. These approaches seem to offer more support for accident investigation because skill transfer requires the additional resources associated with maintaining a pool of investigators. Many reporting systems rely upon the analytical capabilities of one or two investigators [119].

Training programs offer alternative means of both monitoring and supporting the performance of incident investigators. Many reporting systems offer 'refresher' courses. These sessions can be used to introduce new incident and accident analysis techniques [36]. They also provide opportunities to assess and compare investigators' performance in the analysis of case study incidents. Most organisations lack the resources that are necessary to move beyond relatively ad hoc re-training programmes. In contrast, the NTSB is in the process of establishing an Academy for transport accident investigators. This is scheduled to begin operation in April 2002. It will be based in George Washington University adjacent to the U.S. Department of Transportation's National Crash Analysis Center. The intention is that the Academy will build upon existing Investigator Training Courses that are currently only held every six months. The Academy will extend the curriculum and provide a focal point for retraining. It will also provide focussed instruction in the different areas of expertise that are included within the NTSB's multi-party system, described in previous paragraphs. For instance, the reconstructed wreckage of TWA flight 800 will be held at the Academy 'so that future generations of aviation professionals and accident investigators from around the world can learn the lessons that it has to teach' [614]. The Academy will also provide a focus for investigators across the 'international investigative community' [611]. The establishment of this institution forms part of Jim Hall's response to the increasing complexity of many transportation incidents and accidents.

Previous paragraphs have identified a range of techniques that can be used to monitor the performance of incident investigators. Periodic reviews, such as that performed by the HMRI, can identify widespread failures in the analysis of particular failures. Accident investigations can also help to identify more systemic failures. For example, the Cullen report into Ladbroke Grove argued that a 'no blame' culture may paradoxically make staff more likely to accept responsibility for adverse events [194]. It is difficult to see how such reviews might be used to monitor the everyday activity of individual investigators. In contrast, many accident and incident reporting systems rely upon team-based techniques to validate the results of an investigation. The inter-modal exchange of key personnel and the NTSB's 'go team' concepts all explicitly allow for the cross-checking of any analysis before it is released beyond the agency that conducted the investigation. Finally, periodic retraining can be used to ensure that staff are brought 'up to date' with recent developments in investigatory techniques and in specialist areas such as meteorology and structural engineering. The NTSB's new investigator's Academy arguably represents the most significant recent development in this area.

The techniques described above are based around 'traditional' forms of monitoring. Team-based validation, the exchange of personnel, training and retraining have all formed core techniques of Human Resource Management for several decades. There have, however, been a number of more recent initiatives that offer new opportunities to monitor the performance of individual investigators. Many of these techniques are based around the novel computational systems that have been described in Chapter 14. For instance, many incident reports are now stored using relational databases. These systems enable managers to continuously monitor the performance of individual investigators. The same techniques that enable them to identify patterns of failure in a database of incident reports can also be used to identify patterns of analysis or bias in the findings of an individual investigator. In initial trials, we have begun to explore the effects that such information can have upon the overall management of a reporting system. For example, some managers have preconceived ideas about an ideal distribution of causal factors across an incident database. Individual investigators who exhibit a different pattern of analysis are encouraged to look more carefully for those factors that have been neglected in their previous reports. This can lead to potential problems if investigators feel unreasonable pressure to produce a particular pattern of causal findings almost irrespective of the incidents that they have been asked to analyse. Our initial studies have identified further effects, some of which were less easy to predict that the attempts to 'enforce' of normative causal distributions. The intention behind providing an individual with information about the results of their previous investigations is to reduce the problems of frequency and recency bias. There is, however, a danger that this information can have the opposite effect. Showing an investigator that they have identified human error in all recent incidents can reinforce rather than challenge their tendency to identify this cause!

These extensions of existing relational technology are relatively unsophisticated. More recent information retrieval systems offer a number of alternative monitoring tools. Many search engines now routinely construct user models. These models are based around information about an individual's previous retrieval requests. They can be used to make inferences about the user's future information requirements. For example, a frequently used data source may given a higher weighting than one that the user has seldom visited in their previous interactions. User models can also provide insights for the manager of a reporting system. For example, they can be used to determine whether or not an investigator has accessed information about particular aspects of an incident. An investigator's report might rule out human factors as a probable cause even though they have only performed a cursory analysis of the evidence in this area. Conversely, they might focus on particular aspects of an incident to the exclusion of alternative hypotheses.

15.2.3 Intervention Measures

Previous sections have argued that a reporting system can be assessed in terms of the recommendations that it produces. In other words, managers can monitor the effectiveness of the remedial actions that are identified in the aftermath of an adverse event. This information can be used in a number of ways. Previous sections have focussed on the direct benefits that such insights can have upon the operation of safety critical applications. They also fulfill a wider role in helping regulators, the public and government to monitor the operation of a reporting system. A successful scheme should continue to identify areas for improvement. For example, the NTSB issued a document entitled 'We Are All Safer' which stresses the impact that their investigations have had across the transportation industries [604]. They cite safety improvements from more than 150 recommendations in rail passenger car equipment and design, injury reduction and train collision avoidance. These have resulted in seats that are now secured against movement in the event of a collision or derailment. NTSB recommendations are also argued to have encouraged the installation of shatter-proof windows that can also be used as emergency exits. They have caused the installation of overhead luggage racks that have effective retention devices. The NTSB also point to the introduction of passenger emergency briefing cards and too the use of conspicuous levers to help in the operation of doors and emergency windows. The survey also directs the reader's attention towards the less 'visible' effects of their investigations. NTSB recommendations have led to the replacement of railway car construction materials to meet flammability, smoke emission, and toxicity standards. Their analysis has also encouraged the development of new procedures for emergency passenger car evacuation and revised training programs in emergency procedures for service employees. They have helped to introduce mandatory speed and signal compliance checks in certain regions. They have encouraged the use of written notifications to inform employees of speed restrictions and special permission procedures for trains entering out-of-service track sections. NTSB recommendations have also helped to introduce regular crew fitness for duty checks. They summarise the impact of their activities by emphasising the long-term effect of their findings on the industry regulator:

> "As a result of years of rail passenger safety recommendations from the Safety Board, the FRA is enacting regulations regarding passenger equipment safety standards and passenger train emergency preparedness. These regulations will implement many of the recommendations the Safety Board has made to the FRA and the railroad industry to improve the crashworthiness of rail passenger cars and locomotives." [604]

It is possible to criticise the use of such examples to indicate the 'vigour and vitality' of a reporting system. Several of the innovations identified by the NTSB were already being introduced prior to their recommendations being published. It can, therefore, be argued that their intervention helped to expedite changes that were already being made within the industry. This emphasises a point that has been made repeatedly throughout this book. Most incidents and accidents reveal problems that are already well-known by safety managers and other operational staff. Adverse events and near-miss incidents help to focus attention and increase the priorities associated with existing safety concerns. This need not undermine the argument that successful recommendations provide evidence about the health of a reporting system. The NTSB's support for existing initiatives indicates a healthy relationship with regulators and other industry bodies. Equally, there is a concern that any investigatory organisation is independent of such external influences. For instance, the NTSB review describes an incident at Silver Springs when a Maryland commuter train ignored a signal and collided with an Amtrak passenger train. This is the accident described in Chapter 3. The Safety Board found that the crew failed to obey the signals because of multiple distractions and the failure of federal and state regulators to analyse the human factors impact of signal modifications on that rail line. This discussion illustrates two key points. An effective reporting system cannot simply be assessed in terms of the recommendations that it issues. Any monitoring must also account for the way in which those recommendations are received and implemented by operational and regulatory organisations. If all recommendations are accepted then it can be argued that this indicates too close a relationship between the investigators and the recipients of any proposed intervention. Conversely, investigators may propose interventions that are consistently rejected by regulators or more senior safety managers. It might be argued that such situations illustrate the perseverance of an investigator fighting for necessary safety improvements. It can equally be argued that the consistent rejection of proposed recommendations illustrates a break-down in the operation of the system. In either case, serious concerns can be raised about the effectiveness and efficiency of the reporting system.

Mode	Number Issued	Percentage of Total	Acceptance Rate
Aviation	4214	36.2%	82.61%
Highway	1865	16.0%	88.48%
Intermodal	225	1.9%	76.34%
Marine	2234	19.1%	74.75%
Pipeline	1192	10.1%	85.62%
Railroad	1941	16.7%	81.57%
Total	11770	100%	81.99%

Table 15.10: NTSB Safety Recommendations Issued by Mode [615]

Table 15.10 summarises the relative acceptance rates for NTSB interventions in each of the modes of transportation that they are responsible for. As can be seen, the reputation and authority of NTSB recommendations helps to ensure relatively high levels of agreement. However, the difficulty of establishing consensus in the maritime industry, noted in Chapter 13, is arguably again illustrated in this table. It is also important to note the relatively low number of inter-modal recommendations in Table 15.10. This is surprising given that the NTSB is often cited as an example of inter-modal investigation techniques being used to learn lessons that are common to many different industries. There arguments supported both the creation of the ATSB and the US Chemical Safety and Hazard Investigation Board. The relatively low proportion of inter-modal recommendations might be explained by the need to identify a regulatory or industrial organisations to receive such proposals. However, inter-model recommendations can be addresses to more than one recipient. This observation may also reflect the traditional model in which recommendations are focussed towards the particular circumstances of the incident or accident that is being investigated. It remains to be seen whether the inter-modal initiatives being launched by the ATSB will yield a greater proportion of such recommendations than those of the NTSB, summarised in Table 15.10.

Table 15.11 provides a more detailed break-down of the status associated with the NTSB's recommendations following rail-related incidents and accidents. As can be seen, a relatively complex classification system is used to monitor the performance of investigations in this domain. For example, a distinction is made between 'open' and 'closed' recommendations. Open recommendations deal with issues that the NTSB considers still to pose a significant threat to future safety. In contrast, 'closed' recommendations may have been superseded by changes within the industry. Alternatively, they may have been adopted and implemented or they may simply have been rejected by their intended recipients. This detailed breakdown is necessary if recommendations are to be considered as part of the monitoring process. A relatively low acceptance rate can be indicative a large number of different situations. It might suggest a break-down in communication between investigators and industry representations, 'Open–Unacceptable Response'. It can also suggest that recommendations have been amended through successful negotiation, this might be revealed by a relatively high proportion of 'Closed–Acceptable Alternate Actions'.

It is important to emphasise that Table 15.11 provides a very high-level overview of the status of particular recommendations. A reporting system is likely to have a different impact on different aspects of an industry as diverse as the US rail system. For example, many of the innovations that were cited at the start of this section focussed on passenger transportation. More detailed data is required so that safety managers can determine whether recommendations are more likely to be implemented in this area rather than in freight distribution or in rapid transit systems. The NTSB survey does address these different areas by enumerating the particular improvements that have been triggered by their recommendations [604]. For instance, the New York City Transit

Status of Recommendation	Number
Closed–Exceeds Recommended Action	4
Closed–Acceptable Action	1136
Closed–Acceptable Alternate Action	137
Closed–Unacceptable Action	270
Closed–Unacceptable Action/Superseded	14
Closed–Reconsidered	62
Closed–Superseded	19
Closed–No Longer Applicable	110
Total Closed	1752
Open–Acceptable Response	93
Open–Acceptable Alternate Response	2
Open–Unacceptable Response	26
Open–Response Received	15
Open–Await Response	53
Total Open	189
Total Issued	1941
Acceptance Rate	81.57%

Table 15.11: NTSB Railroad Recommendation Status [616]

system has introduced standardisation braking distances and testing procedures. It has also installed speedometers and improved speed control signage. Similarly, a collision involving Greater Cleveland Regional Transit Authority (GCRTA) trains revealed that a train operator had disconnected the automatic cab signal system that provided one form of collision prevention. Coded track circuits were used to transmit speed commands to the on-board train control equipment. To avoid the speed limitation, the operator cut the cab signal to deactivate the control system. As a consequence of the NTSB recommendations, the GCRTA implemented procedures for recording the use of cab-signal cutouts to prevent unauthorised operations.

These specific examples are not supported by more detailed statistics about the frequency of recommendations or the status of those proposals that have already been made. One of the difficulties in this area is that recommendations must be addressed to the many different state and local government organisations. These bodies have the primary responsibility for the safety of the two billion passengers that use Rapid Transit systems each year. Different safety oversight procedures operate in each of these systems. This raises a further more general point; the acceptance of a recommendation by a regulator or other safety body does not imply that the recommendation will be successfully implemented at an operational level. It is, therefore, important that investigatory organisations monitor the effectiveness of their accepted proposals and not simply the overall rate of acceptance, as illustrated in Table 15.11. This is illustrated by the way in which NTSB recommendations argued for mandatory drug and alcohol testing in from the 1970s through to its introduction on US railways in 1986. By monitoring the results of these tests, it was argued that the regulations had helped to reduce substance abuse. Post-accident tests indicated that the number of employees with positive test results fell from 5.5% in 1987 to less than 1% in 1995. Random drug tests showed a similar decline from 1.04% in 1990 to 0.9% in 1995. The success of other recommendations is less easy to establish. For example, the NTSB investigated 29 locomotive derailments in 1991. Diesel fuel spills occurred from ruptured tanks and lines in more than half (56%) of these incidents. The Board issued recommendations that resulted in a joint meeting between the FRA, the Association of American Railroads and locomotive manufacturers. This resulted in a program to collect further data on fuel tank damage and fuel spills. The results of this initiative have taken time to assess because of the delay between revised equipment design and the widespread introduction of these devices across the network. Given the relatively low frequency of adverse events, it took until 1997 before the Board was called upon to investigate two passenger train derailments involving locomotives with the revised 'integral' fuel tanks. The fuel tanks on-board these trains were integrated into their frame structure rather than being suspended within the frame. Integrally tanks also provide higher ground clearance than conventional designs. Investigators concluded that the performance of these enhanced designs 'clearly outperformed frame-suspended fuel tanks' [604]. There was less fuel tank damage and no significant spillage in either of the accidents despite serious track damage. It is, however, difficult to quantify these improvements without considerable additional analysis give that it relies upon a variant of the counterfactual reasoning that has been described in Chapter 11. Investigators are forced to compare the consequences of incidents that might have occurred had the trains been fitted with more conventional fuel tanks.

The arguments cited in previous paragraphs have been drawn from a document that was deliberately intended to promote the investigatory work of the NTSB [604]. It, therefore, provides a relatively positive view of their role in ensuring the safety of complex applications. In contrast, it is possible to identify a more pessimistic or pragmatic view in the technical publications of other investigatory organisations. For instance, the HMRI report into SPAD investigation found that the delays in implementing previous recommendations often meant that new incidents had occurred before previous ones were adequately addressed [352]. Even though one signal at Birmingham New Street had been passed at 'Danger' seven times in eight years, it still took 12 months from the last SPAD incident to install countdown markers. In Railtrack Great Western Zone (RTGWZ), it took ten months to install long hoods on Reading signal R242. The investigators also found strong regional variations in the implementation of recommendations. This led to a meta-level finding that all 'recommendations should be time bound and operating companies should more actively track their completion by setting up their own (monitoring) systems' [352].

It is important to stress that the issues described in this section are generic. They do not simply affect the rail industry. For instance, a recent inspection of UK nuclear facilities also focussed on the efficiency and reliability of monitoring systems that are intended to support the implementation of recommendations from incident reports. They found that some recommendations from incident investigations remained 'incomplete' while others had been expedited as a matter of priority.

> "Although we found information on the state of close-out was being passed to managers, we found little evidence of it being used by managers. Often the only people we could find who were concerned about overdue recommendations were relatively junior staff tasked with keeping the action tracking database up to date. Generally we saw no effective monitoring by managers of those people responsible for closing out recommendations." [640]

Delays and regional variations in the implementation of recommendations are not simply symptomatic of inadequate monitoring. They can also reflect deeper problems, including opposition by both regulators and operators. The following section, therefore, identifies metrics that might be used to monitor the credibility and acceptance of a reporting system.

15.3 Acceptance Measures

The previous section has described ways in which the efficiency of a reporting system can be judged in terms of the recommendations that are implemented in the aftermath of adverse events. There are other ways of reaching similar assessments. For instance, it is possible to monitor the health of a reporting system in terms of those recommendations that are adopted by a regulator but which are violated by their ultimate recipients. The scale of such violations is illustrated by Table 15.12. This summarises the enforcement actions that were initiated by the UK railways Inspectorate between 1996 and 2000. As can be seen the provisional figures for 2000-2001 show a record number of enforcement actions. It might be argued that this illustrates a rising rejection both of the recommendations derived from previous incidents and of the general regulatory framework that supports the UK rail infrastructure. As with previous statistics, however, things are not so straightforward. We have described how a number of high-profile accidents together with structural changes in the infrastructure company have focussed public attention on the safety of the UK railway. It might, therefore, be argued that this rise in enforcement actions is less a reflection of the outright increasing rejection of safety regulations than it is the result of pressure being applied to the Inspectorate to increase their monitoring activities.

	96/97	97/98	98/99	99/00	00/01 (provisional)
Enforcement notices	24	33	21	45	51
Prosecutions heard	6	8	10	11	12
Total fines (£)	233,500	67,500	695,000	1,899,500	1,115,000

Table 15.12: HMRI Railway Enforcement Actions (1996-2000) [334]

This section looks beyond enforcement statistics to identify further metrics that might be used to access the acceptance of a reporting system, not simply the recommendations that it produces. Before looking in detail at these assessment techniques it is important to stress that most reporting systems rely upon the cooperation of many different groups. It, therefore, follows that some of these groups will exhibit different degrees of involvement in a reporting system. For instance, a regulator may offer strong encouragement for the introduction of a system that is opposed by line management in operating companies. Alternatively, Trades Union representatives might support the operation of a reporting system that is not fully supported by regulators. It is also important to emphasise that the public statements of support from some of these groups do not necessarily imply that other, similar organisations will share the same sentiments. For instance, the Transport Salaried Staffs' Association is an independent trade union that represents members in the railway industry, the travel trade, London Underground/Transport for London and London buses as well as road haulage, shipping and ports. They have offered strong support for the CIRAS confidential reporting system mentioned throughout this book. In a recent newsletter they summarise the recent changes that have extended this system throughout the United Kingdom. CIRAS 'is open to Railway Group members and other participating companies, and comprises a core facility supported by three regional centres'. The University of Strathclyde operates the new core facility and runs one regional centre. The other centres are run by a consultancy firm, W.S. Atkins, and a group within the UK's former Defence Evaluation and Research Agency. Railway Safety, an independent company with links to the infrastructure operator, funds the cost of the core facility while the regional centres are paid for by participating companies. The Transport Salaried Staffs' Association note that CIRAS will help employees to report concerns to the regional centres. Centre staff will then conduct follow-up interviews and provide data to a national database. The Association 'supports the important work of CIRAS and is confident that its members in the rail industry will contribute to its effective functioning' [781]. In contrast to this positive message, a number of operating companies expressed initial reluctance to join the scheme. A range of concerns focussed on the usefulness of the data that the system might produce, on the management and confidentiality of the scheme and on cost projections for a national system. This led the Deputy Prime Minister, John Prescott to state that the reporting system would be introduced 'whether or not' the train operating companies wanted it [102]. He argued that CIRAS is "an essential tool to restoring confidence in the industry and getting the actual facts of what is going on". In the aftermath of the Ladbroke Grove accident, however, several industry representatives argued that CIRAS could only provide a short term solution. The need to operate a confidential reporting system was seen as an indictment of the safety culture in the industry because many employees were reluctant to speak openly about safety concerns. A seminar held in preparation for the Cullen enquiry into this accident reached the following conclusions. The final sentence in this statement indicates some of the tensions that can arise between employers and Trades Unions in both the operation of a reporting system and the wider monitoring of safety concerns:

> "There is a problem of finding volunteers to represent the workforce as safety representatives, although this is not universally accepted. Where problems had been encountered,

> it was due to either complacency or employees who were too frightened due to potential victimisation. Trade union representatives can be seen as a nuisance factor and this is an inherent problem of the railway culture. However, it was acknowledged that unions did have a significant part to play in the area of communication, but not at the expense of the normal company communication channels." [466]

Such concerns illustrate the complex and differing attitudes that characterise reactions to voluntary, confidential reporting systems. A similar diversity of opinions can be observed in opinions about mandatory reporting schemes. For example, a review of standard setting across the UK railway found a number of conflicting attitudes towards the value of existing incident and accident investigation procedures [327]. Some of the groups contacted expressed confidence in existing arrangements. For example, the infrastructure company pointed to the introduction of fully independent Chairmen for more important internal inquiries. Others argued that investigations were still conducted with 'insufficient openness'. As a result, other groups such as insurers and consultants had to demand site access. This, in turn, was perceived to have increased concerns over liability rather than focus attention on safety improvements. The different opinions expressed about both mandatory and voluntary reporting systems illustrate the way in which different groups can express different degrees of satisfaction with the same scheme. The following sections argue that these different attitudes can also indirectly influence the effectiveness of many reporting systems. If key groups of workers remain unconvinced about the usefulness and confidentiality of a system, or of regulatory and managerial involvement, then they may be reluctant to participate in its operation. It is, therefore, important to monitoring the acceptance of a reporting system so that such problems can be both detected and addressed before they compromise the effectiveness of a reporting system.

15.3.1 Safety Culture and Safety Climate?

Safety culture forms part of the wider corporate culture that can be used to 'distinguish one organisation from another' [301]. It can be difficult to derive a precise definition of what constitutes a strong safety culture. For instance, Pidgeon and O'Leary identify four different concepts within this term: "responsibility for strategic management; distributed attitudes of care and concern throughout an organisation; appropriate norms and rules for handling hazards and on-going reflection upon safety practice" [681]. Conversely, Reason argues that an organisation embodies rather than possess a safety culture [701]. In other words, the development of a strong safety culture requires 'root and branch' changes to managerial and organisational structures. It cannot simply be grafted onto an existing institution.

Unfortunately, it can be difficult to apply these high-level observations to analyse the particular characteristics of complex, real-world organisations. These problems can be illustrated by a recent review of safety across the Irish railway system. This identified an 'improved safety culture' in the removal of fire risks and in an improved working environment in the underpart of signal cabins. The report also argues that the introduction of elected Safety representatives has also had a positive impact upon safety culture. Such specific improvements can be contrasted with more general observations. It was argued that 'the culture of safety has still not taken root in the staff at ground level' [390]. Workers continue to expose themselves to hazardous track-side conditions with relatively poor protection arrangements. The report also argues that poor morale and the breakdown of management/employee relationships have also had an adverse effect on safety culture. Such generalisations can also be contrasted with specific observations about regional differences in the safety culture within the same organisations. Workers showed greater distrust about managerial attitudes towards safety in the Dublin area than elsewhere in the railway network. This mixture of specific observations and broad generalisations is typical of the types of analysis that are used to support conclusions about the 'safety culture' within particular organisations. Unfortunately, very few studies use the same general or specific observations to support their arguments about safety culture. In consequence, it can be difficult to determine what criteria can best be used to assess the performance of a particular organisation.

In spite of the problems in defining what is meant by the term, many regulators still cite the development of a 'strong safety culture' as a primary aim. This objective is often used to justify the introduction or revision of safety regulations, including the Safety Management System requirements imposed on railway companies by Transport Canada [778]. The maintenance and acceptance of an incident reporting system is often taken to indicate a positive safety culture. There is, however, a need for more detailed metrics to show that regulatory intervention has had the intended effect. The introduction to this section has illustrated the way in which regulators might use the number of prosecutions or enforcement actions as a crude indicator of the 'safety culture' across an industry. A falling number of enforcement actions might be interpreted as evidence that the insights from incident reporting systems are being acted upon without the need for regulatory intervention. A recent survey of attitudes across the UK rail network identified these links between violations and reporting behaviour:

> "If there are rules they should be complied with... A healthy culture would accept the challenge of compliance but would not accept non-compliance. Flagrant disregard of rules needs to be sanctioned in some way. A company should not sanction people for violations that have resulted in an accident if they do not sanction the violations that have not resulted in an accident. There needs to be consistency. For sanctions to be effective the rule that has been broken needs to be seen as legitimate. There is a greater chance of this happening if the employees that operate the systems have a chance to influence and comment on the rules. Therefore if they transgress, they are breaking their own rules. It is important that rules and amendments to the rules are communicated effectively to the workforce, with an explanation of the rule or change. Outlined above are the characteristics of a safety culture moving away from the blame culture and to a more just culture. A just culture will allow for more transparency and for candid reporting, but will not condone reprehensible action." [465]

Others have identified more general links between safety culture and reporting behaviour. For instance, Lucas distinguishes between three different types of safety culture [843]. He argues that the shared perceptions and beliefs that are implicit in these different 'models' can have a profound impact on the types of incidents and accidents that an organisation might experience. Firstly, some organisations exhibit a 'traditional' safety culture. In this approach, the causes of any failure are likely to be attributed to the inattention or carelessness of individual workers. From this it follows that disciplinary actions are the mst likely remedial actions [444]. Alternatively, a 'risk management' approach to safety culture is typified by an engineering view of the human involvement in incidents and accidents. Failures are the results of a failure to correctly design the workers' tasks to their capabilities. Recommendations will focus on changes in operating procedures and on retraining. Finally, Lucas identifies a 'systemic' safety culture. The causes of an incident and accident are analysed in terms of the total working context. In addition to poor task allocation and training, recommendations will focus on mismanagement, on poor communications, low morale, inadequate feedback. The distinctions introduced by Lucas are important because they emphasise the diverse nature of 'safety cultures' within an industry. There are, however, a number of limitations. For example, the idea of a 'systemic' safety culture lacks the clarity of many other 'systemic' approaches to system failure. It is unclear how to measure or even recognise when such an approach has been adequately adopted by an organisation.

This is a significant problem given that it is often necessary to specify some timelimit by which necessary changes should be implemented throughout an organisation; 'as to the length of time within which a company or an industry can see an improvement in their safety culture, many consider that if marked results are not seen within three to five years, then it is likely that the company or industry's approach to developing a good safety culture is flawed' [465]. For instance, the UK Railway Group has devised a safety plan that is intended to take a decade to be implemented based on the premise that 'it may take up to five years for a good safety culture to develop' [465]. It is important that appropriate metrics be identified to help establish when a 'good safety culture' has been achieved. This creates problems because many of the attributes of a safety culture, such as managerial attitudes, cannot easily be assessed or validated. Kjellen, therefore, distinguishes between the abstract notion of a safety culture and the idea of a safety climate, which can be measured [444]. A safety climate denotes 'such aspects of an organisation that are possible to measure by use of a questionnaire-based survey where the results meet statistical criteria for aggregation to the organisational level'. Unfortunately, he acknowledges considerable disagreement over the dimensions that might assess the prevalent safety climate within an organisations. These include:

- management attitudes and commitment;
- involvement of employees in safety management system;
- communication about safety matters between the groups in an organisation;
- risk perception and the attribution of cause in an incident or accident investigation;
- relative priority associated to safety in comparison to other production goals;
- adherence to safety rules and attitudes to the acceptability of rule violations;
- active search for new hazards before incidents take place.

Unfortunately, he also acknowledges that although research has been conducted into safety climate for almost two decades, 'the positive effect of measuring the safety climate for use in feedback to the organisation have yet to be demonstrated' [444]. There are further problems. As we shall see, it can be difficult to demonstrate the reliability of the various instructments that might be used to measure attributes of a safety culture [190]. Surveys that reveal particular attitudes from certain members of staff at particular moments in time do not always achieve the same results when issued to other members of staff or even to the same individuals at different times [395].

The difficulty in assessing the 'safety climate' of an organisation have not dissuaded people from advocating the use of these metrics to monitor incident reporting systems. For instance, the perceived success of the CIRAS system has been cited as evidence of a poor safety culture across the UK rail network. CIRAS supports 'the silent majority who are too scared to report incidents direct to their supervisors and senior management' [466]. Staff are worried about the reaction of their supervisor. They are concerned that they will be disciplined for reporting violationsi. However, it has also been argued that CIRAS can help to correct a deficient safety culture; 'the method of incident investigation is important in developing a proper culture' [466]. This link between the establishment of an incident reporting system and the development of an appropriate safety culture has also been recognised by the US Department of Transport. The Federal Railroad Safety Enhancement Act of 1999 sought 'to reduce human-factor causation of injuries, wrecks, and deaths by improving the safety culture in the railroad industry by expanding and strengthening existing statutory protections for employee whistle blowers' [237]. Statutory protections were extended to cover the reporting of injuries to the railroad, cooperation with an FRA or NTSB safety investigation and refusing to authorise use of equipment, track, or structures that the employee reasonably believes pose an imminent danger to human life. Such initiatives reiterate the expectation that reporting systems should play a positive role in promoting an appropriate safety culture. It, therefore, follows that measurements of the safety climate might trace the impact that a reporting system has upon an organisation. For example, a series of seminars conducted in the aftermath of the Ladbroke Grove and Hatfield accidents found evidence of a 'positive and pro-active' move away from the blame culture and towards a full root cause analysis during investigations. However, they also found that 'indications that frontline staff are not convinced' by initiatives such as the CIRAS Scheme [465]. There was scepticism about whether it was genuinely confidential. This initial concern was reported to have reduced, especially, when the confidentiality of small groups had been protected by analysts generalising the details of a particular incident. There were, however, still problems amongst middle management 'where it was most needed' [466].

The problems of using safety-climate metrics to monitor the impact of a reporting system are illustrated by the different attitudes to the CIRAS system. Some industry analysts that it actually hindered, rather than supported, the development of an appropriate safety culture amongst some workers. For instance, sub-contractors were deliberately excluded from the CIRAS system. Industry surveys revealed the sense of vulnerability and the concerns that sub-contractors felt about an 'us and them' attitude [466]. These concerns were exacerbated by insecure terms and conditions that were offered in response to the increasing financial pressures on the industry. Differences in the safety culture between direct employees and sub-contracting staff were also increased by the introduction of new contractors from the construction industries who lacked specialist railway knowledge. This led to further communication breakdowns, for example in the procedures used to hand-over critical tasks and in the monitoring of safe working hours. These concerns suggest that there are important differences in the safety climate within different sectors of the same industry. Attitudes towards the effectiveness of a reporting system might be very different depending on whether one asked a direct employee or a sub-contractor who was excluded from the scheme.

The overview of UK railways, cited in the previous paragraphs, argues that ‘it is the lowest level of data, such as near misses and non-compliances that do not result in an accident or even adverse effects, that are indicative of the safety culture of an organisation’ [465]. This creates a potential problem. We can use safety climate metrics to monitor the impact that a reporting system has upon the safety culture within an organisation. However, the safety climate within an organisation is assessed by monitoring the submissions to a reporting system. A number of alternative metrics might be used to assess the impact of a reporting system in terms of any changes to a safety culture. For example, previous sections have summarised a broad range of direct measures that include the lost time accident rate or the severity and frequency accident rates. Unfortunately, it can be difficult to agree on and then obtain the information that will be used in this way. Alternatively, safety culture might be assessed by looking at ‘the quality of the relationships within a company and between companies and how effectively they consult and involve their staff’ [465]. Such metrics may ignore the relatively flexible, informal communications channels that are used in smaller working groups across the rail industry [466]. Surveys of staff attitude can also be used to assess the safety culture in an organisation. Transport Canada has produced a checklist to illustrate this approach, based on work by Reason [623]. They emphasise that their safety culture checklist provides no guarantees of immunity from accidents or incidents and that complacency is safety’s ‘worst enemy’. Personnel and managers change so a high score may not be sustained unless the organisation shows constant vigilance. In the following questions, a score between 16 and 20 is indicative of a safety culture that is ‘so healthy as to be barely credible’. Between 11 and 15, the organisation is in good shape. A score between Between 6 and 10 is ‘not at all bad, but there’s still a long way to go’. A result between 1 and 5 indicates that the organisation is very vulnerable. For each of the following questions, a ‘yes’ answer means that ‘this is definitively the case in my organisation’ and adds a score of one to the running total. An answer of ‘do not know’ or ‘maybe’ adds a score of 0.5. Responding ‘no’ or ‘this is definitely not the case in my organisation’ adds zero to the total.

1. “*Mindful of danger*: Top managers are ever mindful of the human organisational factors that can endanger their operations. (Yes/No/Don’t know)

2. *Accept setbacks*: Top management accepts occasional setbacks and nasty surprises as inevitable. They anticipate that staff will make errors and train them to detect and recover from them.

3. *Committed*: Top managers are genuinely committed to aviation safety and provide adequate resources to serve this end.

4. *Regular meetings*: Safety-related issues are considered at high-level meetings on a regular basis, not just after some bad event.

5. *Events reviewed*: Past events are thoroughly reviewed at top-level meetings and the lessons learned are implemented as global reforms rather than local repairs.

6. *Improved defence*: After some mishap, the primary aim of top management is to identify the failed system defences and improve them, rather than to seek to divert responsibility to particular individuals.

7. *Health checks*: Top management adopts a pro-active stance toward safety...

8. *Institutional factors recognised*: Top management recognises that error-provoking institutional factors (under-staffing, inadequate equipment, inexperience, patchy training, bad human-

machine interfaces, etc.) are easier to manage and correct than fleeting psychological states, such as distraction, inattention and forgetfulness.

9. *Data*: It is understood that the effective management of safety, just like any other management process, depends critically on the collection, analysis and dissemination of relevant information.

10. *Vital signs*: Management recognises the necessity of combining reactive outcome data (i.e., the near miss and incident reporting system) with active process information. This involves the regular sampling of a variety of institutional parameters (scheduling, budgeting, fostering, procedures, defences, training, etc.), identifying which of these vital signs are most in need of attention, and then carrying out remedial actions.

11. *Staff attend safety meetings*: Meetings relating to safety are attended by staff from a wide variety of department and levels.

12. *Career boost*: Assignment to a safety-related function (quality or risk management) is seen as a fast-track appointment, not a dead end.

13. *Money vs. safety*: It is appreciated that commercial goals and safety issues can come into conflict. Measures are in place to recognise and resolve such conflicts in an effective and transparent manner.

14. *Reporting encouraged*: Policies are in place to encourage everyone to raise safety-related issues (one of the defining characteristics of a pathological culture is that messengers are 'shot' and whistle blowers dismissed or discredited).

15. *Qualified indemnity*: Policies relating to near miss and incident reporting systems make clear the organisation's stance regarding qualified indemnity against sanctions, confidentiality, and the organisational separation of the data-collecting department from those involved in disciplinary proceedings.

16. *Blame*: It is recognised by all staff that a small proportion of unsafe acts are indeed reckless and warrant sanctions but that the large majority of such acts should not attract punishment...

17. *Non-technical skills*: Line management encourages their staff to acquire the mental (or non-technical) as well as the technical skills necessary to achieve safe and effective performance.

18. *Feedback*: The organisation has in place rapid, useful and intelligible feedback channels to communicate the lessons learned from both the reactive and pro-active safety information systems...

19. *Acknowledge error*: The organisation has the will and the resources to acknowledge its errors, to apologise for them and to reassure the victims (or their relatives) that the lessons learned from such accidents will help to prevent their recurrence." [623]

There is an assumption that this questionnaire will be answered by individual workers reflecting on their experience of the organisations that employ them. This raises interesting issues. One individual can have a very different experience of an organisation than their colleagues within the same team. Individual events can have a profound impact upon answers to general questions such as 'the organisation has the will and the resources to acknowledge its errors...'. In consequence, it may be necessary to aggregate the individual views of many different workers to obtain an overall assessment of the culture within an organisation [863]. Unfortunately, this smooths out the regional and occupational differences that have been noted throughout this section. It is for these reasons that many investigators prefer not to talk about 'safety culture'. The measurement of 'safety climate' and 'corporate culture' raise similar conflicts between the need to generalise and the need to account for differences throughout an organisation.

Previous sections have argued that there can be difficulties in the identification and collection of direct measures for the safety improvements that are attributable to incident reporting systems. For example, a rise in the number of adverse events reported through the system might indicate that the reporting system has failed to deliver necessary safety improvements. Alternatively, it can be argued that the reporting system has helped to increase submission rates or that the incident rate might have been even worse without the reporting system. Further problems complicate the use of measures that focus on the efficiency of the reporting system rather than on safety improvements. For instance, there can often be significant disagreements between an investigator's findings and those of their peers. It can be difficult to distinguish whether such differences arise from the nature of the incidents that they have been asked to investigator or from particular forms of bias that may have affected their causal analysis. The previous section has identified a further set of problems that affect the use of less direct metrics as a means of monitoring incident reporting systems. For example, it is tempting to monitor the impact that a reporting scheme has upon the safety culture within an organisation. Unfortunately, many of the metrics that are used to assess 'safety culture' are themselves derived from the reporting system, such as submission rates. Other measures, such as subjective questionnaires, raise problems because the aggregation of individual returns can hide important cultural differences within a complex organisation.

15.3.2 Probity and Equity

Reason identified three different components of a safety culture: justice, flexibility and learning [701]. Previous sections have argued that flexibility and learning are essential if complex organisations are to respond to the insights provided by incident reporting systems. In contrast, the following paragraphs focus more on the issue of 'justice'. It is important to monitor perceptions about the probity and equity of such schemes in order to assess whether such schemes retain the confidence of potential contributors. Most incident reporting systems depends upon widespread participation in order to ensure that potential insights are not missed through opposition to the scheme itself. Even where there is widespread agreement about the benefits of a proposed system there can be subtle differences of opinion. For example, the General Secretary of the Associated Society of Locomotive Engineers and Firemen wrote to the Deputy Prime Minister in the aftermath of the Ladbroke Grove accident to express his Trades Union's concerns about the future of safety in the railway industry [30] Some of his arguments focussed on the need to increase the involvement of full time officers of the Unions in cross company safety meetings. This would support the exchange of safety information and would encourage 'an open safety culture'.

The General Secretary's views are both important and influential because they represent informed opinion and carry political weight within the industry. The survey based techniques that can be used to aggregate different attitudes towards safety culture often fail to account for such strategically important opinions. In contrast, safety managers must carefully consider the political weight of such views if they are to ensure participation in a reporting system. For example, the General Secretary went on to identify Union concerns over the punitive nature of many investigations. This results in a 'secretive culture' that stifles information sharing about safety issues. In contrast, he argues that the Society's representatives should be involved in setting up a 'no blame' policy for driver retraining following adverse events. His response reiterates the Union's support for the CIRAS initiative. The crucial point to consider here is, however, that this support is offered in the context of these wider safety concerns about secrecy, punitive investigations and the lack of consultation. Again, such issues cannot easily be extracted from aggregate responses to high-level questionnaires about 'safety culture'. In particular, the Union response cites reports from the infrastructure operator that some operating companies have blocking the extension of CIRAS because they fear vindictive employees would abuse the system. He argues that the opposition from operating companies also stems from a concern that spurious submission will waste the time of their managerial staff. The General Secretary argues that these concerns reflect a negative attitude that is based on outdated prejudices. It is 'a sad reflection on management within the railway industry'. Such comments illustrate the recursive nature of many safety concerns. Not only do they reveal the attitudes of the person writing the letter, they also reveal their attitudes towards the opinions that they believe others hold about a reporting system. It is difficult to construct direct question that might elicit such information. The following sections, therefore, describe alternative qualitative techniques that might be used to monitor particular attitudes towards the 'probity and equity' of a reporting system.

Informal interviews and focus groups can provide insights into the fears that particular individuals have about voluntary and confidential reporting schemes. These same techniques can also be used to expose the attitudes of managers and regulators that might not be so readily obtained from more formal questionnaires and surveys. The recent independent review of Australian rail safety provides an example of the way in which these techniques can provide important insights into attitudes towards incident reporting [55]. The review deliberately canvassed a wide range of opinions; 'the industry is diverse, and it was expected that different organisations would have varying views'. The study consulted major track managers; all major freight and passenger operators and a sample of the smaller passenger and freight operators; all rail safety accreditation authorities; the Rail Safety Committee of Australia; representatives of new entrants, rail client groups, workers safety committees; rail heritage and tourist groups; the Industry Reference Group Chair as well as Commonwealth and State agencies. They note the willing of these groups to discuss safety issues 'forthrightly and at length' and observe that there 'were many common themes, with differences often a result of the particular circumstances of the organisation concerned'. It is interesting to note that this review avoids any attempts to define or characterise a single 'safety culture' across the Australian rail network.

The survey was deliberately intended to solicit views rather than derive numeric values. Hence the monitoring can be seen as qualitative rather than quantitative. This makes it particularly important to justify the use of particular elicitation techniques and then to validate any subsequent interpretation of the information that is received. Unfortunately, the published accounts of the review provide summary information. Relatively little information is provided about the elicitation process and conclusions are often presented without reference to the supporting data that was obtained from the parties mentioned in the previous paragraphs. For example, the independent report argues that"there is a level of co-operation which is being achieved in specific instances between the industry participants and parties such as the coroner and the occupational health and safety authorities". This includes the sharing of evidence and interview results in the aftermath of incidents and accidents. Such observations are instructive because they illustrate important strengths in some regions. However, the report does not describe the reasons why or how this level of co-operation was achieved in particular locations. Such details, arguably, lay outside the scope of the review. The high-level nature of these comments may also reflect the need to protect the confidentiality of the contributors. It does, however, illustrate the way in which qualitative reviews can lack the grounding provided by the statistical analysis of more direct monitoring techniques.

The potential need for additional evidence is also evident when the report identifies problems with the existing arrangements for incident and accident investigation. It reports a residual concern 'that specific competencies are needed to assess the cause of railway incidents and this presents a danger that evidence needed to determine the operational or technical causes may be lost, or recommendations which compromise railway operational best practice may be imposed' [55]. Such observations illustrate how the results of qualitative surveys can be used to support the monitoring of incident and accident reporting systems. Unfortunately, additional details are required if regulators and managers are to address these high level criticisms. The independent review does provide some details in a subsequent analysis of the 1996 Intergovernmental Agreement on National Rail Safety between the Commonwealth, States and the Northern Territory. The various parties expressed concerns about many aspects of incident and accident investigation:

- there was a lack of clear protocols for the parties on site and for containment of the site after an adverse event;

- there were concerns about the independence of investigations undertaken by the regulator, operators, managers and other 'interested' parties;

- there was a perception that restrictions had been placed on the means by which the results of an investigation were communicated to the industry. There had been a failure to alert the industry of potential 'hot spots';

- there was concern over the different focus of investigations between those to which 'no-blame' was attached and those aimed at prosecutions. Participants were worried about the potential for manipulation of the self-incrimination provisions of rail safety legislation;

- there was a perception that undue delays had occurred in finalising investigations or making cause information available where litigation was expected.

These are valuable observations. For example, they indicate that investigation protocols should be drafted and then publicised to increase confidence in the results of any analysis. Assurances should also be given about the independence of investigations. Unfortunately, it can be difficult to prioritise these concerns so that management and regulators can prioritise their allocation of resources. Relatively little information is provided about how whether these concerns were shared across the national system or whether they were isolated within particular geographical regions and functional groups. Some information of this nature is provided. For example, 'most (of the respondents) argue for at least a minimum national role' in the coordination of safety management activities across the rail system. Similarly, the report states that 'most respondents' recommended that this minimum role should include the collation and analysis of statistical data on incidents and accidents, feedback on the causes of accidents through safety bulletins, the coordination of major incident and accident investigations, ensuring the 'standardisation of interpretation'. The report also identifies dissenting views from this majority opinion. National operators were concerned that proposals for a national organisation have not addressed the problems of multiple jurisdictions, of inconsistent analysis and of fees to support investigatory organisations. They also argued that most rail activity focussed on intrastate business and commuter functions. It, therefore, made little sense to focus so much attention on a national body. In addition to the dissenting views of national operators, the report also identified the concerns of 'many small operators'. Some of the these companies are isolated from the mainline network and, therefore, did not consider that national regulation should affect their operations.

This report, therefore, illustrates both the strengths and weaknesses of qualitative approaches to the monitoring of incident and accident reporting systems. Surveys can be used to sample the diverse views that help to form the different safety cultures within complex organisations. Unfortunately, the lack of empirical data can make it difficult to assess whether or not a particular opinion is shared by the majority within a particular group. Confidentiality agreements can also prevent analysts from providing access to the tapes and notes that support particular conclusions. They can also isolate the reader from the contextual information that might help to interpret particular comments. Many of the findings of these enquiries can appear to be based on supposition rather that the precise statistical findings of more direct techniques. Qualitative techniques can, however, also be used to summarise a broad range of opinions that might otherwise have been hidden within particular metrics. For instance, the previous paragraph identifies important differences between national operators, small scale companies and most of the remaining groups that were surveyed.

15.3.3 Financial Support

Previous paragraphs have described how qualitative techniques yield important insights about the fears that particular groups hold about reporting systems. Workers express concerns about the probity and equity of investigations that are coordinated by management on behalf of regulatory and investigatory organisations. Managers and operating companies are worried about the undue influence of national regulatory bodies. They are also concerned that employees will waste finite managerial resources by generating spurious reports. This latter argument introduces a further means of monitoring the performance of a reporting system. It can be argued that the funding arrangements, which support a scheme, can provide valuable insights into the perceived success of the system. If companies provide financial security without without regulatory obligation then it might be argued that the scheme is well respected. Conversely, continual funding reviews and a lack of invest might be interpreted as important signs that a reporting system is failing to provide valuable insights into necessary safety improvements.

This line of argument is supported by attempts to justify continued public investment in the work of the NTSB. In 1996, $850 million was allocated to the FRA to support a regulatory rail safety program. $39 million dollars was allocated to the NTSB to ensure the safety of all forms of transportation, including oversight of the Federal Railroad Administration. To place this in perspective, the Federal Transit Administration received $4 billion to fund rail transportation infrastructure and equipment purchases. The NTSB recognise the 'substantial investment' in rail transportation safety and, therefore, acknowledge their responsibility to act as 'the eyes and ears of the American people at accident sites' [300]. Similar comments can be made about the FRA's responsibilities to elicit information about 'less-serious' incidents and near miss events.

It is possible to identify a range of different funding mechanisms that have been used to support both incident reporting systems and the wider regulatory infrastructure that supports them. For example, the UK Civil Aviation Authority is unusual because it was established with the aim that the industry should pay the cost of its own regulation. Income is partly derived from from licensing charges. Such charging schemes can lead to inequalities. For example large operators may claim that they are subsiding smaller companies if relatively more regulatory time is spent on their concerns. Conversely, smaller companies might claim that they subsidise larger operators if a unit fee is charged irrespective of the size of an organisation. These comments have also been made about the levies that support reporting schemes [444]. In order to address these concerns, the CAA also raises income from a levy on airlines that is based on on passenger kilometres. In contrast to the UK CAA, the costs of US aviation regulation are recovered from transportation users by various indirect taxation. For example, through levies placed on passenger tickets. This can lead to further problems. For example, such charges have been criticised because of the impact that they can have upon particular forms of transport. For example, Australian rail operators have pointed to the relative subsidies that road transport operators receive in relation to rail operating companies. Road operators do not meet the full costs of maintaining the national network while rail has to pay track access fees. The imposition of further overheads to support enhanced safety regulation, including the national incident reporting systems mentioned in previous sections, exacerbates this perceived imbalance [857].

The funding of incident and accident reporting on Australian railways is more complex than the previous paragraph suggests. Each operator pays a safety accreditation fee that varies between the States even though there are mutual recognition agreements. It is, therefore, possible for an operator to seek accreditation in a State that is different from the one in which they conduct the majority of their business. The inconsistencies in funding also reflect deeper variations in the safety regulations that are enforced in different areas of Australia. Previous sections have described how these include the regulations covering the reporting of adverse events. It can, therefore, be argued that funding mechanisms might provide useful metrics to help monitor incident reporting systems. A recent series of reports urged that '... the Commonwealth takes a strategic approach to provide consistency in rail safety standards and practices for the national track' [857]. It was also recommended that 'a single annual fee for accreditation should be payable only in the jurisdiction of principal activity'.

Both license-based and taxation-based models of funding create financial pressures during a 'down-turn' in the economy. Transportation companies are, typically, faced with falling revenues. They must, however, continue to meet the licensing costs that are necessary to support incident reporting systems and the wider regulatory framework. In this model, financial burdens remain on the operators. In contrast, under a taxation based scheme, both the regulator and the operator are hit hard by falling sales. This creates particular problems for incident investigators. It can take well over a year to train an analyst [197]. Skilled and experienced staff cannot easily be dismissed in response to short-term market fluctuations.

There are alternative funding mechanisms. As mentioned previously, both the UK HMRI and the Health and Safety Executive impose specific charges for the work that they conduct. Time is invoiced for each quarter or half an hour spent on an investigation. This leads to invoices that contain several thousand entries and fee recovery takes between 3 and 4% of HSE/HMRI resources [467]. It can be argued that this represent an inefficient use of scarce resources. In particular, if this model were used on a subsidised national railway then public money would simply be transferred from rail operations to rail regulation. UK railways, therefore, operate a levy scheme to support the

Office of the Rail Regulator. This overseas the economic aspects of market intervention. A similar scheme is proposed for safety regulation. This has the strong advantages of simplicity and economy in contrast to other forms of funding [467]. As we have seen, however, it raises important questions about the scale of the levy to be placed on each individual operator.

It is clearly important that those who pay for reporting systems should realise benefits that are in proportion to their investment. This creates potential conflict because there is no direct relationship between funding and control in the area of safety regulation. For instance, the International Civil Aviation Organisation (ICAO) require that the investigation of accidents and serious incidents is conducted by an independent organisation. This principle is also reflected in recent European directives on the regulation of the aviation industry (such as 94/56/EC). This distinction is not embodied within the UK rail industry where Railway Safety is a not-for-profit, wholly-owned subsidiary of the infrastructure company, Railtrack Group PLC. Railway Safety does, however, operate under a separate management structure from its sister company, Railtrack PLC, which is responsible for operating the infrastructure under the Railtrack Group. It is unclear whether such a situation could continue if the European Commission implements the proposed extension of independent requirements to other modes of transport [467].

It can be argued that a 'healthy' incident reporting system should have the same financial and operational independence as investigatory organisations within civil aviation. The rules that separate accident investigation bodies from other regulatory or commercial organisations do not extend to incident reporting systems. Most are financially dependent on the agencies that implement their recommendations. For instance, the CIRAS system was initially funded by the rail companies that operated in the region that it covered. Such close relationships can create concerns; investigators may be reluctant to propose recommendations that are unpopular with financial contributors. In consequence, a National Steering Group was established to oversee the national CIRAS system. The members of the steering group include individuals from Railtrack Safety and Standards, Railtrack Line, Railway trade unions, the Association of Train Operating and Freight Operating Companies, the Infrastructure Safety Liaison Group and an independent human factors specialist [196]. A Charitable Trust has also been created to 'promote and protect the independence and integrity' of the CIRAS system. Again the members of this trust include a representative from Railtrack Safety and Standards, Railway trade unions, a human factors academic, a member of the Rail Passenger Council, a representative of the core facility service provider, and representatives of rail employers.

By monitoring the level of funding that is made available to a reporting system, it is possible to assess the investment that companies are willing to make in these schemes. As we have seen, however, economic trends can reduce the financial support that is made available to a reporting system. The previous section has also argued that additional managerial devices must be used to ensure the independence of many schemes, especially if they receive high levels of financial support from regulatory and commercial organisations. A number of further problems complicate the use of financial metrics to assess the health of a reporting system. Incident and accident investigation require specialist skills. It can be difficult to recruit and retain necessary staff. One recent survey argued that there were no independent rail incident investigators anywhere in the UK; 'consultants who do not work for Railtrack do no exist' [467]. The lack of independent investigators is compounded by structural and organisational problems that act as barriers to recruitment even when funding exists. For example, railways are often perceived to lack the 'glamour' of other high-technology industries. This creates problems in recruiting the best graduate, technical skills. The difficulties of staff recruitment and retention are compounded by the government Civil Service pay structures that operate within the UK HMRI. When there is competition for scarce talents 'the HMRI has been limited in what it could do by a lack of good people to take work forwards' [467]. It is important to stress that these recruitment problems also affect investigatory agencies across a broad spectrum of industries, including mining, nuclear and off-shore oil production, and in many different countries not just the UK railways.

15.4 Monitoring Techniques

The previous pages in this chapter have introduced broad distinctions between the different techniques that might be used to monitor the success or failure of an incident reporting system. Particular attention has been paid to the problems of interpreting the information that is provided by many of these monitoring techniques. For instance, an increase in the financial resources that are allocated to a reporting system may not be sufficient to attract skilled personnel. Conversely, a fall in regulatory contributions can increase the independence of some reporting schemes [467]. In contrast, the remainder of the chapter focuses in more detail on a subset of these monitoring techniques. Brevity prevents a complete exposition, however, the intention is to summarise the issues that must be considered before investing in a particular approach to the validation of a reporting system.

It is important to emphasise that the particular techniques used to audit a reporting system will depend upon the scale of the scheme and the organisation that it is intended to support. This point is reiterate by Transport Canada's guidelines for the development of railway Safety Management Systems [780]. They argue that monitoring and audit frequencies should depend on the size of the railway, the risks involved in their operations and the previous safety performance of the organisation.

> "Larger railway companies will likely have the staff and expertise necessary to establish auditing processes and teams, although they may choose to hire external resources to obtain specific skills or assistance. Smaller companies that may not have the resources to conduct an audit program internally may be able to obtain assistance from a variety of sources, including senior railways with which they interchange, consultants and professional auditors." [780]

Some authors have argued for the continuous monitoring of the performance of incident reporting systems, for instance using the direct measures introduced in previous sections [444]. For small scale systems, this can divert critical resources away from the analysis of adverse events. It may, therefore, only be possible to conduct periodic monitoring every six or twelve months [119]. Fortunately, a range of computer-based monitoring systems can be used to reduce the costs and hence increase the frequency of monitoring activities. The costs associated with some monitoring techniques, such as observational analyses, can dissuade safety managers from exploiting these techniques even on larger-scale schemes. The following sections, therefore, use previous applications of these techniques to provide an impression of their relative costs and benefits for the monitoring of reporting systems.

15.4.1 Public Hearings, Focus Groups, Working Parties and Standing Committees

Many different types of meeting can be called to help monitor an incident reporting system. Most of these hearings are called in the aftermath of particular failures. They, therefore, typically considered reporting systems within the context of a wider safety management system. It is rare for public hearings, focus and working groups or standing committees to concentrate exclusively on the utility of a particular scheme. This broader focus does not, however, prevent these meetings from providing important insights about the performance of a reporting system. For example, many focus groups begin by looking at the perceived causes of a particular incident and then go on to question the reasons why lessons had not been learned from previous, similar incidents. The following sections, therefore, briefly describe the ways in which these different venues can be used to provide feedback about reporting systems.

Public hearings provide a means of assessing general attitudes towards incident and accident reporting systems. These meetings are often called to review general safety concerns in the aftermath of major failures. For instance, the FRA held a series of public hearings following a number of incidents in which passengers had been unable to escape from trains in the aftermath of a derailment or collision [234]. The catalyst for these meetings was the Silver Springs incident described above; a Maryland commuter train ignored a signal and collided with an Amtrak passenger train. Such public meetings pose a considerable challenge to those who must both organise and chair them. There is a danger that pressure groups will attempt to promote their views and exclude those of other groups with valid concerns. Equally, however, it is important that the convenor of a meeting should not be seen to stifle debate by imposing a rigid control over the proceedings. The FRA have well-rehearsed mechanisms for addressing these potential problems. The dates of a proposed public meeting are published in the Federal Register. Members of the public must then notify a clerk of their intention to speak. They must also submit three copies of their planned oral statement by a date that is specified in the call for participation. Members of the public are notified that their submission has been received by the FRA. Their written submissions are then made available for examination by other potential participants and by representatives of 'interested parties' prior to the meeting. This procedure has several merits. Firstly, it alerts the meeting chair to potential conflict. Secondly, it helps to ensure that any questions of fact can be raised and resolved before the meeting so that any subsequent debate can be based on reliable information.

Public meetings are often held to identify concerns that have not been addressed by working groups, focus groups and standing committees. For example, the FRA's public hearings were called in response to an interim report that was published by a working group on Passenger Train Emergency Preparedness. It is difficult to establish clear distinctions between these other forms of meeting. The terms 'working group', 'focus group' and 'standing committee' are often used synonymously by both regulators and operating companies. In general terms, however, a focus group can be thought of as an informal meeting that is held to consider a particular series of issues. The meeting need not arrive at a particular plan of action but may produce broad recommendations about the items being discussed. In contrast, a working group can be thought of as a more formal device to both consider particular issues and then act to resolve them. The life time of the working group usually ends with the successful resolution of the items being considered or by the implementation of their recommendations. A standing committee, typically, has greater longevity. They are often intended to provide a continuing point of reference for the consideration of long-standing issues. All three of these devices can and have been used to monitor the success or failure of incident reporting systems.

As mentioned, public meetings often attract participants that have a particular perspective of, or vested interest in, the issues that are being discussed. For example, passenger groups, environmental protection organisations, the proponents of road transport have all actively participated in recent public meetings on rail safety [336]. Such organisations are well placed to represent particular views within the wider community. They may not, however, reflect the diversity of attitudes held by the general public. In consequence, many organisations rely upon focus groups to investigate perceptions about the safety performance of particular industries. These meetings have the benefit that participants can be selected to deliberately reflect a broad cross-section of views. For example, the FRA used focus groups to assess compliance with railway operating rules. The intention was to assess whether corporate culture had an influence on potential violations [245]. This study illustrates how focus groups play a particularly important role in analysing the causes of common failures. As we have seen, incident reports can often provide information about what happened. It is far more difficult to understand why particular patterns of failure occur across an industry. The FRA in using this technique have sought to provide additional analytical information than that which is normally provided through their mandatory reports scheme.

Focus groups can be used to directly assess particular attitudes towards the operation of an incident reporting system. For instance, the US Bureau of Transportation Statistics undertook a series of workshops to identify 'stakeholder' concerns about the reliability and accuracy of accident and incident information [116]. Their concerns should not be surprising, they reiterate concerns that have been raised throughout this book. The participants drew attention to data quality. They were concerned about both the under-reporting and the over-reporting of particular types of adverse events. They were worried by the lack of uniformity in completing reports. They voiced concerns over exclusions that removed reporting requirements from some transportation workers. Typical comments include 'there needs to be better information and it needs to be of a higher quality', 'there needs to be better data on results', 'accuracy is a challenge because of budgetary problems and different interests' and 'it is difficult to get accurate, undiluted information on human error and performance' [116]. The focus groups also revealed concerns over the relevancy of data produced by the Bureau of Transportation Statistics. Industry participants were concerned to ensure that the right information was being collected and that data that was duplicative or no longer useful was not collected. This final observation is highly instructive. Focus groups are one of the few mechanisms that can be used to obtain feedback about the overheads that imposed on potential contributors by reporting requirements. Many of the other measures, such as submission rates or intervention metrics, take little account of the costs that a system might impose upon potential contributors.

Focus groups are more commonly used to discuss concerns that arise in the aftermath of high-profile accidents and incidents. Many of these concerns centre on the failure of reporting systems to prevent the occurrence, or mitigate the consequences, of the adverse event. This can be illustrated by a recent seminar held in the aftermath of the Ladbroke Grove accident. A focus group explicitly considered the role of incident reporting as part of a wider review of employee attitudes to rail safety [466]. This seminar included present and former railway staff, signalers, Control Room Operators, incident and accident investigators and project managers. All participants appeared in a personal capacity, however, and were not intended to ast on behalf of any particular organisations. A list of the questions were circulated to the participants before the meeting. They were asked to send in brief comments that were then circulated to the other members of the focus group before the meeting.

1. How concerned about safety are those who work on the railways?
2. What are the main concerns with respect to safety on the railways?
3. How important do those who work on the railways consider safety to be, relative to other issues such as punctuality and reliability of train services?
4. In practice, are safety requirements compromised by commercial considerations?
5. Has the fragmentation of the rail industry had an adverse effect on safety? If so, in what respects and for what reasons?
6. Is there uncertainty or confusion as to who is responsible for what with respect to safety on the railways?
7. On a personal level, are those who work on the railways aware of their duties and responsibilities with respect to health and safety issues?
8. Are unsafe acts and conditions tolerated on the railways? If so, do they go unreported? How can this problem, if it exists, be addressed?
9. Is there a mechanism whereby those who work on the railways can express safety concerns to those in positions of authority within their organisation? Is the mechanism effective? Are their concerns addressed and acted upon?
10. Is the confidential incident-reporting system on the railways used? Is it trusted and respected by the workforce? Is it effective?
11. How are safety issues communicated from directors and other policy-makers to the workforce? Are safety issues given enough emphasis? Are there sufficient safety-related incentives?
12. How often do those who work on the railways receive visits and/or safety briefings from supervisors and senior managers? Are the briefings effective?
13. How often do those who work on the railways receive formal training on safety-related issues? Is the training effective?
14. Does the reliance on contractors and sub-contractors for track repair and maintenance prejudice safety?
15. What is the delegates' understanding of the safety case regime? Have those who work on the railways seen their company's railway safety case or 'assurance case'? Is this a document they use or on which they rely? How does the document relate in practice to more prescriptive requirements such as the Rule Book?

16. What should be the appropriate balance between the use of broad objectives on the one hand, and detailed prescriptive rules on the other, to achieve safety on the railways?

17. What can be done, or should be done, to improve safety on the railways?

This example illustrates several important features about the use of focus groups to monitor incident reporting systems. The seminar was was not intended to reach particular conclusions on any of the questions. The intention was review employee perspectives on safety in the rail industry. This reveals an irony in the use of the term 'focus group'. These meetings frequently move from a focussed set of issues to more general and wide ranging discussions. It is, therefore, important that the facilitator or organiser retained control over any meeting without dictating the content of the discussion. Our case study meeting initially focussed on what the employees' main safety concerns, including issues of leadership, responsibility and accountability. Only then did the focus group concentrate on communication mechanisms, including incident-reporting. The meeting also focussed on many other issues ranging from the employment of contractors to training and the use of UK railway's rule book [466].

As mentioned, focus groups are often used to provide general feedback about attitudes towards a reporting system. For example, the UK meeting described an initial scepticism about whether CIRAS 'provided a genuinely confidential reporting scheme' [466]. The members of the focus group were found that 'experience of 34 months working the system showed that it was excellent, a lot had been learned and that there was no breach of confidentiality'. In contrast to focus groups, working parties are typically expected to provide detailed recommendations. For instance, the Health and Safety Executive recently established a working group to 'deal with the problem' of vehicles crashing onto railway lines from overhead bridges [113]. This group collated evidence and analysis from a large number of incident and accident investigations. Their analysis of this collated evidence recommended more barriers, improving the road layout and introducing better signs for drivers.

Working groups often coordinate their activities with those of other, broader forms of consultation. For example, previous paragraphs mentioned the public meetings that were called following reports from the FRA's Passenger Train Emergency Preparedness Working Group. It was argued that the FRA must become more proactive in order to minimize the consequences of future accidents; 'even minor incidents could easily develop into life-threatening events if they are not addressed in a timely and effective manner' [234]. The establishment of this working group might be seen as an implied criticism of existing incident and accident reporting systems. In this case, accidents such as the Silver Springs collision have demonstrated that more action needs to be taken to mitigate the consequences of any future failures. In this view, the working group is established to supplement systems that have failed to adequately address existing safety problems. Equally, however, it can be argued that the establishment of the working group illustrates the success of existing reporting systems. The need to consider emergency preparations has been established from the analysis of previous incidents.

The participants in a working group are typically chosen to ensure that a broad range of interests are represented. They, therefore, play an important role in assessing the feasibility of the recommendations that are produced from a reporting system. The expertise and experience of the participants can often help to identify implementation concerns that were not initial recognised by incident and accident investigators. For example, one recommendation from previous collisions and derailments was that the FRA should require the introduction of booklets and videotapes to illustrate equipment and describe entry and evacuation procedures on trains. The Working Group pointed out that 'that pilferage of on-board emergency equipment is a serious problem on many passenger railroads, and that specifically focusing the attention of passengers on where the equipment is located would only exacerbate the problem' [234]. They also argued that frequent travellers probably already knew where the equipment was located and would not, therefore, benefit from such additional information. This case study provides further examples of the way in which a Working Group can provide feedback on the recommendations derived from previous incident reports. For example, Amtrak used the meetings to point out the difficulties of introducing emergency preparedness booklets and videos across its network. Not only would they have to distribute this information on many thousands of rail services, they would also have to send them to emergency responders throughout the

United States. Subsequent mailings would also have to be used to ensure that any information was up to date. The FRA considered these comments to the Working Group and invited commentators to 'suggest either how Amtrak can best comply with the emergency responder liaison requirement as set forth in the proposed rule, or whether the final rule should establish a different standard for railroads that operate in territories with large numbers of potential emergency responders to contact' [234].

The FRA working group illustrates the use of such meetings to assess the recommendations that have been produced in response to previous incident reports. It illustrates the use of these meetings to look at particular safety issues, in this case emergency preparation on passenger trains. Similar techniques have been used at a higher level to review incident reporting systems within the wider context of a national regulatory framework. For example, the Ladbroke Grove and Southall rail accidents led to an industry-wide review of safety management of UK railways. A working group was established under the Department of Transport, Local Government and the Regions. This conducted a review of the Safety and Standards Directorate within the infrastructure company, Railtrack. As with the FRA case study, the findings of the working group were informed by and helped to inform public inquiries. In this case, the Department of Transport working group implemented a number of significant changes pending the recommendations of the Public Inquiry into the Ladbroke Grove accident. The scale of these changes cannot be underestimated. The working group initiated the transfer of responsibility for determining whether or not another train company was safe to operate from Railtrack to the Health and Safety Executive [688]. Railtrack's Safety and Standards Directorate were transformed into a separate, non-profit making company with an independent board of directors within the Railtrack Group. The objectives of this new organisation were to provide 'safety leadership' to the industry and take a more dynamic approach to setting and updating standards. More significantly given the focus of this book, the new Railway Safety body was to 'establish a more effective regime for safety audit, incident investigation and ensuring that corrective action from audit and investigation is taken' [688]. Previous paragraphs have described the Precursor Indicator Model (PIM) that has been developed by Railway Safety to support this more pro-active approach to safety audits. The Working Group's general review of rail safety, therefore, triggered changes that 'revolutionised' both the operation and monitoring of mandatory incident reporting across the UK rail network [691].

Regulatory and governmental agencies are responsible for commissioning most working groups. Professional organisations, industrial bodies and pressure groups have also starting investigations into the success or failure of incident reporting systems. For instance, the UK Royal Aeronautical Society's Human Factors group has established a Rail and Aviation Working group [375]. This aims to share human factors expertise, resources and best practice from the aviation communittee with representatives of the rail industry. This Working Group was explicitly established with the Royal Aeronautical Society because as 'an impartial professional charity' it can provide the intellectual resources and unbiased refereeing that may not be available from other similar bodies. Representatives are drawn from Railtrack, the Rail Industry Training Council, the Aviation Training Association, British Airways, the UK Flight Safety Committee and individual train operators amongst others. This Working Group has focussed on transferring lessons from the operation of aviation reporting schemes, in particular British Airways's BASISindexBASIS [658], into the emerging national rail systems. The objectives and even the existence of such a group provides important insights into the perceived health of existing reporting systems within the UK rail industry. The perceived need to transfer skills and techniques from the aviation domain into the railway industry implies a relatively low regard for existing rail systems. The working group description concedes that the aviation safety record is not perfect. However, there is no recognition that techniques might be propagated back from the railway domain to support aerospace safety management.

Public hearings, focus groups and working groups all tend to have a limited duration. Public hearings and focus groups are called to identify particular concerns on topical issues. Frequently, they are used to gather feedback about the management of safety in an industry following high-profile failures. For instance, they may provide insights into the reasons why reporting systems fail to prevent an adverse event. Working groups are similar in that their longevity is bounded by the publication and implementation of recommendations. In contrast, standing committees provide a common point of reference for long-standing concerns. They can be used to coordinate the work of focus groups, of public hearings and of working groups. For instance, the Royal Aeronautical Society's Human Factor's group has standing committees on crew resource management, on maintenance engineering and on air traffic management amongst other topics. These groups are intended to monitor developments, set up 'focus teams' and advise the main committee on specific Human Factors issues. The Human Factors group can itself also be seen as a standing committee; it coordinates the Rail and Aviation Working group mentioned in the previous paragraph.

The UK Railway Industry Advisory Committee provides an example of a government sponsored standing committee. It was established by the Health and Safety Commission in 1978. The Railway Industry Advisory Committee 'plays an important role in providing a consultative forum where all interests within the railway industry can meet and reach consensus on how to progress health and safety proposals and other related developments within the industry' [336]. Meetings are chaired by the Chief Inspector of Railways. Seven employers' representatives and 'balanced' by a similar number of employees' representatives who are nominated by railway trade unions. Passengers and the general public are represented by two members from the Rail Passengers Council. The membership of the committee has been reviewed and revised several times to reflect the changing structure of the industry since privatisation. This process is an important difference with the other feedback meetings mentioned in this section. The limited longevity of working groups, focus groups and public hearings makes membership changes less important than they are for standing committees. As mentioned, standing committees often coordinate the work of these other groups. The Railway Industry Advisory Committee supports a Freight Sub Group; an Occupational Health Working Group; a Prevention of Trespass and Vandalism Working Group; a Research Working Group and a Human Factors Working Group. Each of its working groups have terms of reference and plans of work that are approved by the main Committee and their Chairs report to the main RIAC Committee. Each of these groups draws upon incident and accident reports as an important means of informing their activities. For example, the Occupational Health Working Group has used analyses of previous injuries to draft of industry-specific guidance on manual handling for employers and employees on the railways. The Research Working Group has started two initiatives on track-side safety and on the effects of safety messages on influencing the behaviour of railway passengers. Each of these activities was motivated, in part, by their interpretation of recent safety statistics derived from the various industry reporting systems. The members of the Human Factors working group helped to promote the CIRAS scheme as a means of identifying the safety concerns of operators [318]. The Railway Industry Advisory Committee working groups also helped to monitor the use of the RAVERS fault tracking and reporting system following the Southall accident. This was seen as a short term solution to a situation in which most operating companies had computerised facilities to log faults and produce trend reports but 'the ability to share data nationally is being compromised by industry moves to 'stand alone' systems' [318]. It is important to note that the Railway Industry Advisory Committee blurs some of the distinctions that were introduced in previous sections. There is no suggestion that the working groups will be suspended once the human factors or maintenance issues have been satisfactorily 'resolved'! It might, therefore, be better to refer to these groups as sub-committees that will continue to support the work of the standing committee. The key point is, however, that these bodies provide many different industry stakeholders with the ability to address particular issues over a prolonged period of time. They are not simply established to address the findings of a particular investigation. It is also important to note that incident reporting systems provide a vital information resource to the members of these committees. It should not be surprising, therefore, that the Railway Industry Advisory Committee's working groups have addressed the development of various fault monitoring systems and confidential reporting schemes.

Public hearings, focus groups, standing committees and working groups provide valuable information about attitudes and opinions about particular reporting systems. They, therefore, often provide post hoc information in the sense that opinions are often formed in the aftermath of adverse events. They identify concerns without necessarily offering clear guidance about constructive solutions. There are, of course, exceptions to these generalisations. It is important, however, that safety managers and regulators have access to alternative techniques that can be used to assess the utility of particular reporting systems. Incident sampling techniques address this requirement; representative subsets of previous failures can be examined to determine whether a range of alternative techniques might have yielded further insights into the causes of adverse events.

15.4.2 Incident Sampling

The term 'incident sampling' covers a number of different techniques that extract a sample of reports that have been submitted about adverse events. These techniques differ in the criteria that are used to choose a particular subset of events. For instance, analysts may attempt to extract a random sample. Alternatively, they may base their selection on incidents from a particular functional subsystem or geographical area. Incident sampling can also be focussed on particular levels of severity. For example, monitoring activities may concentrate on those failures that had the greatest potential adverse consequences. Once the subset has been identified, each incident is analysed to assess the quality of the causal analysis, to validate any potential recommendations and so on.

A recent 'Assessment of Investigations into Signals Passed at Danger' on UK railways illustrates this approach [744]. This investigation was conducted by W.S. Atkins in order to monitor the efficacy of revised investigation regulations following the Ladbroke Gove accident. These revisions required that HMRI inspectors investigate each major SPAD incident in addition to any enquiry conducted by the railway companies involved. The Atkins report was partly intended to compare the results of the HMRI investigations with those of the rail operators. They were also intended to interpret their findings in the light of the HMRI's own separate analysis of the railway companys' investigation techniques [349]. The Atkins report examined the conclusions reached from each of these different enquiries in order to identify the 'value added' by having HMRI inspectors perform their own independent analysis of each high-severity SPAD in addition to company investigations. Of the 146 SPADS investigated by both the HMRI and the companies, 13 were selected for further analysis by the Atkins report. It is difficult to identify the precise criteria that were used to inform the selection of this subset. However, a six stage methodology was used to guide the monitoring process.

1. *Data collection.*
   The project began by reviewing the 146 incidents investigated by both the HMRI and industry investigators since October 1999. The results were collated to compare the root causes identified by one or the other or both investigations. The analysis also attempted to identify root causes that might have been overlooked in both previous investigations.

2. *Review company investigations.*
   The analysis then used the collated data to identify causal patterns that might not have been identified by the previous investigations. The Atkins report attempted 'within the limits of the relatively small sample' to identify differences in the effectiveness of investigations between different regional zones. This review also provided insights about the consistency and thoroughness of company investigations compared to those of the HMRI.

3. *Review HMRI investigations.*
   The HMRI reports were also critically reviewed to identify the relative strengths and weaknesses of their analysis. As mentioned, the intention was to identify ways in which the HMRI investigations might 'add value' to company reporting procedures.

4. *Independent analysis.*
   A small sample of 13 SPADs were then analysed in greater detail to identify any root causes not identified by either the HMRI or the company reports. This was intended to determine whether these investigations considered 'the full list of possible causes' and then proposed 'suitable measures' to prevent any recurrence [744]. This stage of the analysis was intended to uncover any correlation between the root causes that were missed and the type or category of SPAD's being investigated. It was also hoped that this analysis might improve our understanding of the reasons why any root causes might have been overlooked.

5. *Collation of root cause data*
   The results of each of the previous stages were collated to summarise all of the root causes identified since October 1999. This data then informed a series of more detailed statistical analyses to identify trends and patters across the SPAD incidents.

6. *Proposals and recommendations*
   The final report identified the strengths, weaknesses, trends and Zone differences observed in both investigation processes. Proposals were also made for ways in which the benefits of the subsequent HMRI investigation might be achieved without the need for a second investigation. Finally, the Atkins report proposed selection criteria that can be used by the HMRI to 'confirm that the Industry investigations are achieving consistency in depth of analysis and conclusions reached' [744].

The results of this process showed considerable agreement between the HMRI and the company investigations. However, the report's writers stress that their comparison focusses on 'the quality of the investigations rather than the results and recommendations' [744]. This is significant because it might be argued that the report, therefore, overlooks the benefits of more direct intervention following the HMRI report. The comparison did, however, illustrate some of the problems that arise when comparing different reporting systems. For example, the independent review concluded that although there were differences in the root causes found by the HMRI and the company enquiries, both were 'equally valid'

The third stage of the methodology, described above, was based around a subjective comparison of the quality of the investigations conducted by the HMRI and the railway companys. The score 3 represented a 'robust report', 2 was assigned for a 'good report' and 1 was used if the report was poor and had 'significant shortcomings in either analysis or conclusions' [744]. Of 228 investigations, only 10 (4.4%) were subjectively classified as 'poor'. No incidents were received a score of 1 for both investigations. In 93 SPADS, the HMRI score was equal or worse than the company score. In 21 (18.4%) incidents the HMRI were assessed as 'adding value' to the industry investigation. The analysis also identified weaknesses in both the HMRI and the company investigations. These can be summarised as follows:

- *Incident Selection.*
  The criteria that were used to select the incidents that were to be subject to both company and HMRI investigations was skewed towards shunting incidents. These were argued to be of relatively 'low consequence' and their over-representation was perceived to indicate an imbalance in the SPAD severity classification scheme that informed the selection process.

- *Terminology.*
  Company and HMRI investigations used different terminology. For example, HMRI and Railway Safety reports used SASSPAD for 'Starting Against Signal SPAD' while in other contexts the same acronym was taken to mean 'Start Away From Station SPAD'.

- *Special Exercises.*
  Special exercises, for example top gather information about hand-signalers, increase awareness of particular safety issues and help to elicit reports about certain classes of adverse events. Atkins' review argued that 'care must be taken' to ensure these initiatives do not compromise 'mainstream' SPAD investigations.

- *National Issues.*
  Although the HMRI had identified issues across many different regional operating zones, some issues had 'slipped through the net in spite of clear evidence of a trend being available' [744]. These included situations where the driver had been forced to apply power even though the signal was at red.

- *Balance between Human Factors and Infrastructure Issues.*
  It was argued that in both sets of investigations, there was a tendency to see human 'error' as

a cause without exploring further into the infrastructure issues that may have made the error more likely.

- *Organisational Issues.*
  The review concluded that both the HMRI and the company reports tended to under-emphasise organisation issues unless they were of an extremely serious nature. They found that 'it is only on rare occasions that we have seen focussed recommendations addressing supervisory or management weakness' [744].

- *Follow-up of Issues Raised.*
  Atkins expressed a concern that in some cases the HMRI had identified significant problems without initiating follow-up actions. They admitted that this concern was not, however, supported by direct evidence.

- *Miscommunication.*
  The review found that incidents involving communication failures were typically blamed on the driver even though other personnel, including signalers, may have contributed to the adverse event.

- *Aspect Sequence.*
  Both the company and the HMRI reports failed to give sufficient consideration to the aspects of the signals that the drivers had encountered immediately before the SPAD incident. In some industry reports 'the cautionary aspects which might have had a significant bearing on the incident have been completely ignored' [744].

- *Technical Complexity.*
  The HMRI occasionally misunderstood or oversimplified the issues involved in complex investigations. In other cases, they used inappropriate language or obfuscated the issues in a way that created an unhelpful image of the Inspectorate.

- *Compliance and Effectiveness.*
  The report argued that the HMRI should be more forceful in challenging previous industry practices especially when it was obvious that remedial action might be costly. They argued that 'lack of adverse comment in certain cases might be construed as acquiescence' [744].

This list illustrates the range and number of insights that can be drawn from relatively focussed monitoring techniques. It should be emphasised, however, that these observations were based on the subjective analysis of the Atkins staff. Similarly, their work was not supported by a formal methodology that might have supported the monitoring of other reporting systems. Such caveats need not, however, undermine the importance of their findings. For example, they were able to identify fundamental differences between the ways in which different company's interpreted causal information. East Anglia operated two different classification systems. In some reports, immediate causes were distinguished from underlying causes. In other reports, a three-tier hierarchy distinguished immediate causes from basic causes and root causes. Great Western A two-level hierarchy: Immediate Cause; Underlying Cause London North Eastern either operated a three-level hierarchy involving immediate, basic and root causes or a complex two-level hierarchy. In this approach, immediate causes were distinguished from 'underlying causes/personal factors or underlying causes/job factors'. Southern region operated an alternate two-tier hierarchy that distinguished conclusions from underlying causes or a three-level hierarchy involving immediate, basic and root causes. This led the Atkins review to suggest a simplified structure distinguishing three different levels. The first level describes primary events and special circumstantial factors. At the second level, basic causes are identified. These include human factors and ergonomic issues. They also include infrastructure problems and procedures or instructions. The final level documents the underlying organisational, managerial and supervisory causes. Such proposals and the diverse approaches operated within each zone indicate the importance of monitoring the operation of similar incident reporting systems across national industries. Experience within the aviation industry has shown that these different classification schemes act as significant barriers to the exchange of important safety-related information [308].

The Atkins review of the SPAD investigation processes illustrates the monitoring of reporting at two levels. Firstly, the report is itself an attempt to monitor the integration of, and value added by, the duplicate company and HMRI investigations. Secondly, the HMRI is itself an indirect means of monitoring the company investigation procedures. Hence there is a sense in which the Atkins report monitors the monitoring system. This can be seen in the review of the criteria that the HMRI used to determine which SPADS to investigate. The Atkins report chose the sample of 228 incidents because these adverse events were the ones where the HMRI had already chosen to 'duplicate' the company investigation. The report raised a number of concerns about the criteria that the HMRI used to select their sample for further investigation. Atkins described some of the incidents that the HMRI monitored as 'low risk, low value'. These included SPADs in which the trains may have travelled relatively large distances beyond the red signal but at relatively low speed and with little risk, for instance within a depot. In many of these cases, the HMRI were able to add little value to the company reports.

The report into the HMRI monitoring also raised more general questions about the limitations of incident sampling techniques. Investigating a sample of all SPADs incidents 'inevitably casts doubt on any statistical data' and especially 'the extraction of cause data'. The Atkins investigators were concerned that a widespread causal factor might go unnoticed if the resulting SPADs do not meet the selection criteria for the HMRI monitoring. They cite the example of several incidents in which the SPADs follow shortly after a train starts on a yellow signal. This type of incident is likely to result in a short distance SPAD as the signal changes to red. Hence, they will not be investigated by the HMRI. The Atkins team suspect that these incidents may be much more prevalent than either HMRI or the industry currently believe. Thee is a certain irony in the sampling criticisms made in this review. Many of these adverse comments might also be levied at this meta-level review because part of their analysis depends on a subset of the HMRI sample. Leaving aside this caveat, the Atkins report goes on to identify a revised set of criteria that might be used to guide the HMRI's decision to launch an investigations alongside a company enquiry. For instance, the HMRI should investigate an overrun which results in injuries or fatalities to either passengers or staff. They should also investigate an overrun which results in damage to the infrastructure, or damage to traction units or rolling stock. These detailed criteria are extended to include broader categories such as incidents that meet criteria defined in occasional special studies and 'a random selection of SPADs'. These requirements are intended to address the problems that can arise when particular definitions inadvertently exclude incidents involving certain causal factors, such as those described above.

In addition to suggesting alternative conditions that might be used to guide the HMRI sampling of SPAD incidents, the Atkins report also describes clear objectives for any future monitoring by the HMRI. The main aim of this activity should not be to 'duplicate' industry investigation. Instead, the HMRI should establish what happened, identify the implications of what happened and determine the effectiveness of any proposed recommendations. The HMRI need to know what happened in order to evaluate the risks of any repetition. They need to know the the implications of what happened in order to assess the potential consequences of any recurrence. They need to assess the effectiveness of any recommendations to ensure that industry proposals will address all the 'components' of the SPAD. In particular, the Atkins report identifies situations in which disciplinary action has been taken against the driver while infrastructure weaknesses were overlooked. They conclude their review of the existing monitoring functions by observing that:

> "Our investigations have revealed a wide variation in the quality of root cause analysis being undertaken. Used properly and intelligently it is a powerful tool for extracting all the implications from an incident. Used mechanistically, and we have seen a number of examples of this, it can lead to some root causes which, whilst they may expose valid weaknesses, bear little relevance to the SPAD and, when corrected, will have negligible effect in preventing a recurrence. Too often these irrelevant issues are being pursued at the expense of more serious ones which are being ignored." [744]

As mentioned, a range of different criteria can be used to identify the sample incidents that can be used to monitor many reporting systems. The Atkins study drew on the sample of SPAD incidents that were investigated by both the HMRI and by company investigators. The HMRI, in turn, used a number of complex criteria to determine which of the company investigations they would also look into. For example, one aspect of the criteria focussed on the length of the overshoot that resulted from the SPAD. The Atkins report was not the only example of meta-level monitoring to be triggered by the Ladbroke Grove accident. This also influenced the HSE to serve two 'Improvement Notices' on Railtrack, the infrastructure provider. These required that they produce a plan, with fixed implementation dates, to reduce the risk of a future SPAD at the 22 signals with the worst safety record across the network.

The signals were chosen because they have a record of multiple SPAD incidents. SPADS are divided into four categories. Category A incidents occur when a train passes a signal at danger without authority, other than those SPADs defined in Categories B, C and D. Category B SPADs occur when a train passes a signal at danger without authority because a stop aspect or indication was not displayed with sufficient time for the driver to stop safely at the signal. Category C SPADs are occasions where a train passes a signal at danger without authority because a stop aspect or indication was not displayed with sufficient time for the driver to stop safely at the signal, because it was returned to danger in an emergency in compliance with rules and regulations. Category D SPADs are those occasions when vehicles without any traction unit attached, or any train which is unattended runs away past a signal at danger [356]. The 22 signals involved in the review were the site of Category A SPADs. These incidents are also classified according to eight severity ratings. A level one SPAD is the least severe with an overrun of no more than 25 yards and no damage or casualties. Severity rating eight is the worst and is used for incidents resulting in fatalities. The review of Railtrack's actions was concerned with 'the identification, understanding and mitigation of the causation factors increasing the likelihood of signals being passed at danger' [356]. The consequences of SPADs at each signal was, therefore, of secondary importance. This contrasts with the previous Atkins report that focusses more on the high-consequence incident that were investigated both by the HMRI and by individual companys.

The HSE enforcement notices were based partly on a statistical analysis of previous SPAD incidents and partly also on a recognition that previous recommendations had failed to prevent the recurrence of adverse events. The statistical analysis showed that for signals that had been passed at danger three or more times, there was only a 4% probability that the SPADs are entirely due to random causes [356]. Where there were four or more SPADs, the probability becomes close to zero. The results of this analysis triggered an enquiry to determine whether investigators had identified all of the infrastructure elements or environmental factors that make SPADs more likely to occur at these signals. The enquiry was jointly conducted by the HMRI and by W.S. Atkins. The focus of the investigation was on Railtrack's response to Improvement Notice I/RJS/991007/2 requiring action to mitigate the risk of signals being passed at danger at the 22 signals that had been passed at danger most often between 1990 and 1998. The previous incident reports for each of the signals were examined. Discussions were held with Railtrack Headquarters and with each of Railtrack's seven Zones. Meetings were also held with representatives of Train Operating Companies and some infrastructure maintenance contractors. All of the signals were viewed from a train cab and discussed with drivers. Reviews were also held with HMRI and Atkins representatives. The methodology used in this investigation is instructive because the report describes further constraints on their terms of reference that affected the manner in which they examined incidents involving the 22 signals. For instance, the HMRI was already conducting a separate inspection of the operating company's systems for driver management in order to assess whether they were adequately addressing the causes of human 'error'. The review, therefore, focussed more on the infrastructure issues that were under Railtrack's direct control than the driver-related factors that were largely the responsibility of the operating companies. A further two of the signals, SN63 and SN109, on the exit from Paddington Station were excluded from the review. They were close to the site of the Ladbroke Grove accident and hence were covered by a separate remedial plan. The meta-level review of Railtrack's actions on the 22 signals also revealed some of the dangers of monitoring incident reporting systems. For example, some of the signals were again passed 'at danger' after the investigation was completed and Railtrack's remedial actions had been approved by the HMRI. The report addressed potential criticism of their monitoring by arguing that 'further improvements may only come from improved management of driver competence and fitness, but Railtrack is also investigating whether there are further improvements which can be made to the infrastructure' [356].

The monitoring activity for each signal was intended to determine whether Railtrack had taken effective steps to understand the probable cause of previous SPAD incidents. The HMRI/Atkins review resulted in requests for Railtrack to carry out further risk assessments on some signals. Two of the signals required remedial actions that could not be completed before the review was published. It was concluded that the remaining signals had received adequate attention from Railtrack. Sufficient attempts had been made to identify likely causal factors and to apply appropriate measures to reduce risk. The report argued that 'in most cases it was also clearly evident that there had been close co-operation with the Train Operating Companies both in identifying and applying the risk mitigation and in ensuring that drivers were well briefed' [356]. Railtrack and the operating companies also revealed 'encouraging' evidence of an improved understanding about the factors that were likely to increase the risk of signals being passed at danger and of measures that might mitigate those factors. These more positive remarks were balanced against a small number of exceptions. At Manchester Piccadilly, a blanket speed restriction had been implemented on HMRI advice. This reduced the risk of future SPADs but the underlying causes of the original incidents had not been adequately identified. Other incidents had taken a significant amount of time to address so that effective remedial action were not implemented until long after the original incident had been reported.

This section has described how a range of different sampling techniques can be used to focus monitoring activities on particular types of incidents. Resources can be concentrated on incidents that receive particular attention within the reporting system, such as the SPADs that were investigated by both the HMRI and individual companys. Alternatively risk-based criteria might be used, including the severity assessments that informed the selection of the 22 worst SPAD signals. The Atkins review of the HMRI and company investigation processes also identified the limitations of sampling techniques. The selection of particular incidents can systematically exclude other incidents. If an investigation were extended to include these other adverse events then the monitoring might identify potential weaknesses in the underlying reporting syste, It is for this reason that the Atkins report was careful to propose detailed conditions that might trigger HMRI investigations in addition to the initial company investigation. A number of further limitations affect incident sampling techniques. In particular, there is little point identifying detailed sampling criteria if insufficient incidents are being submitted in the first place or if it is clear that systematic biases affect the reports that are being contributed about adverse events. These caveats are not a significant problem for SPADs where it is highly likely that any incident will be noticed by signalers and other railway staff. In other domain, however, these issues encourage regulators and safety managers to seek alternative means of monitoring their reporting systems.

15.4.3 Sentinel systems

Monitoring is not simply used to assess the health of a particular reporting system. It can also help managers assessing the potential strengths and weaknesses of alternative reporting techniques. For example, there must be some means for comparing the results of different causal analysis techniques or of different form designs. The importance of such comparisons is illustrated by the Safety Case requirements that govern UK railways [350]. These documents are intended to persuade regulatory authorities of the adequacy of an operator's safety management plan. The Safety Case must provide evidence to show that they will:

- " encourage staff and others to report accidents, incidents and other events that have or could have affected health and safety;
- ensure fair and equitable treatment of those whose actions are examined as a result of an investigation;
- investigate incidents and accidents (including potential or actual instances of ill health) to determine immediate and root causes;

- provide investigators with suitable competence, seniority and authority to undertake impartial effective investigations;
- provide resources to investigate, categorise and implement action to prevent repetition;
- match the investigation response to its potential severity;
- co-operate effectively with other duty holders to ensure that all lessons are shared, learned from and acted on;
- review findings from investigations (both internal and external) periodically to ensure that technical and managerial inadequacies are corrected;
- feed back to staff and others the results of investigations; and provide procedures for all aspects of an investigation." [350]

At present, most safety cases simply assert that particular steps will be taken in the event of an incident or accident. As we have seen, however, the HMRI, Lord Cullen and W.S. Atkins have criticised the adequacy of existing reporting systems. Monitoring techniques provide a means of addressing these criticisms by providing evidence that a proposed approaches provides greater benefits than the potential alternatives. It is important to stress, however, that such comparisons should not jeopardise the operation of an existing reporting system. There is a clear concern that any short-term trials might lose confidence in a successful scheme or lose data about failures that might later prove to be essential in preventing future accidents. Sentinel systems provide a means of obtaining pre-hoc information about a reporting system without forcing changes throughout large-scale reporting schemes.

Chapter 14 briefly described the main features of Sentinel reporting systems [262]. Rather than running a national or regional reporting system, Sentinel schemes elicit information from a small sample of 'representative' units. This approach has numerous benefits. For example, it avoids the costs with larger scale national systems. Sentinel systems also focus training and 'awareness raising' resources on the selected units so that participation rates can be raised above those normally associated with mass schemes. Mass reporting systems are often referred to as 'passive' because they do not actively encourage each member of staff to contribute to the system. In many cases, these higher levels of participation would overwhelm the analytical resources of the scheme. In contrast, smaller-scale Sentinel applications are referred to as 'active' because they seek to encourage closer participation in many different aspects of the reporting process. The smaller scale of Sentinel systems also enables additional resources to support the causal analysis and identification of recommendations from the smaller number of adverse events that are identified by Sentinel systems. It can, therefore, be argued that these schemes provide more reliable insights that mass reporting systems. There are also some limitations. Unless the sample institutions are carefully chosen then it is likely that some incidents will go unreported. It is for this reason that Sentinel systems are often used to complement rather than replace larger scale reporting schemes. This parallel approach has obvious benefits for the comparison of different techniques. The continued operation of a mass system enables safety managers to collect incident data using established techniques. The introduction of a limited number of small scale trials enables comparisons to be made with these existing approaches.

A number of problems complicate the use of Sentinel systems to guide the evaluation of alternative reporting techniques. The additional resources that are typically associated with these schemes can prevent accurate comparisons being made with the lower levels of investment that are possible in mass systems. Sentinel systems elicit more information than mass schemes almost irrespective f the particular techniques that are being used. There are also problems associated with longitudinal trials. Sentinel systems usually involve the introduction of new techniques. The novelty factor involved in learning to implement these innovations can increase motivation and involvement beyond the levels that are associated with the routine operation of mass reporting systems. Any observed advantages might decline if the techniques used in a Sentinel system were extended throughout a national scheme over a prolonged period of time. These objections limit the conclusions that can be drawn from 'direct' comparisons between the results obtained by Sentinel systems and those provided by larger-scale systems. In contrast, these techniques have been used to help validate the results obtained by more passive approaches. For example, the distribution of information obtained from a Sentinel system can be compared to assess the coverage of a national scheme. Sentinel systems can be used to uncover incidents that are under-represented in larger schemes. Hence the focussed application of active reporting techniques help to validate rather than directly compare particular incident reporting systems. Alternatively, different approaches can be trialed in separate Sentinel systems. For example, additional resources might be used to promote the application of PRISMA within one company while another is encouraged to use Tripod. Again, however, direct comparisons can be complicated by the different operating characteristics and safety records of the firms that are involved in the study.

It has been difficult to find any well-documented report of the application of Sentinel reporting within the railway industry. Pasquini, Rizzo and Save have used a variant of this approach to support the analysis of SPADs on Italian railways [664]. In this case, they specifically developed a reconstruction and causal analysis technique that was intended to support the investigation of these incidents. The first stage involved the production of video recordings during physical reconstructions at the site of the SPAD. Focus groups then discussed the film together with relevant technical document and testamonies. These discussions were then used to generate a matrix diagram similar to the MES diagrams introduced in Chapter 11. The actors in this case included the train driver and an on-board signal repeating system. This safety device is similar to the Automated Warning System described in Chapter 5. The validation of the investigation technique involved the cooperation and training of an investigative team, including two drivers. Instead of analysing recent SPAD incidents, as would have been the case in a full Sentinel trial, the comparison of the new techniques with existing approaches was based on a post hoc analysis of three SPADs on Italian railways between April 1997 and November 1998. In consequence, their study focussed on differences in the analysis of these incidents rather than on any improvements in the elicitation of information about adverse events. In particular, they argued that the new methodology helped to identify latent problems with the way in which the signal repeating system was operated. Drivers saw the warnings as a nuisance that were to be dismissed as soon as possible rather than as valuable safety information. These findings were contrasted with the insights provided by the existing reporting systems which focussed more on inattention and lack of concentration.

As mentioned, there are few examples of Sentinel systems being used to support the analysis of railway reporting schemes. Arguably the best documented example is provided by a research project that was funded by the FDA . At the start of this study, the intention was that the Sentinel system would act as a supplement to a mass reporting system. Towards the end of the study, the costs of operating the national system led many involved in the evaluation to argue that this approach could replace the existing scheme. The study address many of the problems mentioned in previous paragraphs by recruiting 23 different facilities for a twelve month period. They secured the support of Study Coordinators who were, typically, risk managers for hospitals and directors of nursing for nursing homes. These individuals participated in orientation and training sessions. These covered the purpose of the Sentinel system, project background and goals, comparison of voluntary and mandatory reporting procedures, project plans and confidentiality procedures. The initial work on the project identified 'large gaps in the knowledge of facility clinical staff regarding the importance of reporting adverse medical device events' [262]. Participating facilities were also provided with video materials for clinical staff training. These encouraged staff to follow their facility's internal procedures for reporting of adverse events. The investigators also contacted the Study Coordinators in each of the participating facilities to gather information about each facility's reporting procedures. Information about these procedures was made available to staff after the videos, mentioned above, had been screened.

The architecture used in the Sentinel evaluation involved Study Coordinators sending incident information to a central group of analysts. The analysts then telephoned the Coordinator to acknowledge receipt of the submission and to confirm any additional information that may have been obtained in the interval after the submission. Coordinators were also encouraged to contact centre staff with more general questions related to their work as Risk Managers. They requested the names of contacts at the FDA, specific information about device tracking regulations, how to use software for filing reports and also whether they were required to report certain events to FDA and manufacturers. Once a report was received by the project staff, they were reviewed with a nurse and a specialist in medical informatics. This preliminary analysis was used to determine whether follow-up requests were required to elicit further information about the adverse event.

The Sentinel project received 315 reports from 23 units between October 1997 and November 1998. 286 reports were submitted through the post in special envelopes provided to the study, 3 were sent by fax and 26 were reported by telephone The telephone reports were particularly instructive because Study Coordinators could tell the analysts about particular problems that they had experienced in completing the paper-based forms. The investigators argued that 'there is reason to believe that the level of DEVICENET reporting activity was far above the average for hospitals in the MEDWATCH system' [262]. It was estimated that if the 5,500 hospitals in the US national reporting system contributed at the same level as the Sentinel facilities then they would receive more than 100,000 contributions each year. The actual total for 'health care facilities' in 1998 was only 5,000. All of the submissions came from hospitals. More than half of all reports came from one large hospital, and a second large hospital contributed another 15% of the reports. It is instructive that even though additional resources were focussed on the participating institutions there were no reports from the six nursing homes. This was explained by the observation that nursing homes are extremely tightly regulated. There was, therefore, strong management concern that negative information might come to the attention of authorities.

As mentioned, the relatively small scale of Sentinel systems enables safety managers to encourage submissions about events that might otherwise overwhelm a national system. In the FDA study, it was determined that only 14% percent of all reports described events that could to have been submitted under the existing mandatory schemes. 56% described events that fell under the Sentinel system's voluntary guidelines. The remaining 30% fell into a gray area; 'it was not clear whether they were mandatory or voluntary' [262]. It was argued, however, that few of these events would have otherwise been reported. Many of these reports related to incidents in which it was difficult to assess whether the patient had suffered a serious injury. The additional resources devoted to the Sentinel system also supported a number of analyses that are not normally performed in larger scale reporting systems. Two senior nurse-analysts reviewed all of the submissions and classified them as ery urgent, urgent, routine monitoring or well-known problem, or not important. Approximately one-third of the reports (113) were assessed as being urgent or somewhat urgent. However, only 19 of these incidents were clearly assessed as falling within the mandatory reporting system. About half of the mandatory events (51%) needed only routine monitoring. For example, incidents were assigned to this group if any problems were already well-documented and if the regulator had already taken action to address them. In contrast, 30% of the 175 voluntary reports were rated as very urgent (2) or somewhat urgent (50). Of the 95 reports that did not directly fall under either the voluntary or mandatory regulations, 44% were either very urgent (2) or somewhat urgent (40). The Sentinel study also examined the way in which Study Coordinators classified each incident. The results of this analysis were significant because these classification represent the primary means of pattern matching in the mass reporting system, for example using the automated retrieval tools described in Chapter 14. The investigators felt that about a third and a quarter of the codes were incorrect [262].

Sentinel based reporting systems are not a panacea. The lack of submissions from nursing homes illustrates that this approach cannot guarantee the participation of all potential user groups. There are further concerns. For instance, the types of facilities that are recruited to many Sentinel studies may already exhibit a high degree of awareness about safety-related issues. If this is not the case at the start of the study then the additional resources that are allocated to the promotion of health and safety can quickly alter the behaviours of many of the operators and work groups that participate in the study. Hence the types of incident information that is provided by a Sentinel system will rapidly become atypical of the adverse events that affect the rest of the user communittee. This can be interpreted as a variant of the Hawthorne effect introduced in Chapter 5. Users will alter their normal working behaviour if they know that their behaviour is being directly or indirectly monitored. A number of potential solutions have been proposed for this problem. In particular, observational techniques can be used to identify particular behaviours that may support or weaken the operation of an incident reporting system.

15.4.4 Observational Studies

As we have seen, Sentinel systems focus additional resources to support incident reporting in a small number of 'select' institutions. This very process of selection and the additional support can help to ensure that the sample facilities no longer resemble other units within the same industry. Hence the information that they provide many not be representative of the adverse events that occur at other sites. It can also be difficult to interpret the insights derived from focus groups and interviews. Operators can express views that are not reflected in their subsequent behaviour. For example, they may strongly support the operation of a voluntary incident reporting system but fail to contribute to a scheme even when they witness an adverse event. Techniques that rely upon the statistical analysis of incident reports suffer from similar limitations. It can be difficult to identify the reasons why particular types of incidents are not reported or why certain groups of operators are reluctant to participate. The following paragraphs describe how workplace studies and other observational techniques from the field of sociology can be used to address these criticisms.

There have been many notable attempts to use techniques from the field of sociology to provide insights into the working lives of railway staff. McKenna has investigated the strategies that railway personnel have used to maintain their standard of living during times when the railways were contracting [531]. Salaman looks at the way in which drivers' attitudes changes towards their occupation and their colleagues [721]. There have also been studies of union responses to changes in management structure [63]. Heath, Hindmarsh and Luff, However, point out that relatively few of these studies focus on the everyday working activities of railway personnel [339]. There are some exceptions. For instance, Gamst has conducted a detailed study of the work and attitudes of US locomotive engineers [284]. Even this study has, however, focussed on workers' opinions and preoccupations rather than the manner in which they accomplish their everyday tasks. There has, however, recently been a move towards applying many of these sociological techniques to provide more direct insights into working behaviour. Heath and his co-authors are amongst the leading figures in this area. Others include John Hughes and his colleagues in air traffic management [374] and Berg in the field of medical safety [78].

In contrast to many previous design techniques, these studies do not focus narrowly on the operation of high-technology systems. In contrast, they consider human-human as well as human-machine interaction. There is also a concern to consider the way in which diverse communication media, including physical artifacts such as pencil and paper, are used to coordinate and inform group activities. Many of the proponents of this approach have written about 'rescuing' the study of technology from cognitive science which concentrates too narrowly on the psychological characteristics of individual users. In practical terms, these 'workplace studies' involve participants joining the groups that they are observing for prolonged periods. They will often follow the same shift patterns as the individuals that they are studying. This is important because it helps the observer to build up a mass of background information that may be necessary to understand the significance of the events that they witness. It can also provide some indication of the prolonged impact that stress, fatigue and other workplace factors can have upon operators and managers.

During these periods of observation, investigators compile field-notes. They can also use audio and video recordings. Clearly, however, the conspicuous compilation of these records can remind workers that their actions are being observed. The nature of these records depends partly on the context in which the study is taking place and partly also on the forms of analysis that will be used after information has been elicited. For example, conversation analysis provides important insights from studying the vocabulary and structure of workers' conversations. This technique is only feasible if transcripts can be reconstructed from field-notes or other recordings. Other forms of analysis require less direct records, such as the construction of social networks to model the way in which different working groups interact [693]. They may, however, require longer periods of observation to ground any potential conclusions in observed behaviour.

There are clear ethical problems in exploiting these techniques to monitor incident reporting systems. For example, observers may witness adverse events and 'near misses' that are not notified by any of the operators who were involved. Other observers have seen users struggled to operate computer equipment that they themselves were familiar with. This creates a considerable dilemma in many safety-critical contexts. If the observer decides not to act then there can be adverse consequences. Conversely, an ill-advised decision to intervene can exacerbate rather than resolve a potential incident.

Brevity prevents a complete introduction to the many different approaches that have been developed to support observational and workplace studies. In passing, however, it is important to stress that most of these technique specifically avoid the generation of hypotheses before the study is conducted. Such concepts should emerge during the observation as more information is gathered about the workers and the context of their daily lives. This guiding principle helps to ensure that analysts do not selectively filter their observations to support pre-formed hypotheses. It is also important to stress that some 'ethnographers' deliberately reject the idea that observational techniques should be used to support particular theories [303]. This argument stems from the discussion that was introduced in Chapter 11. Causal asymmetries complicate the task of explaining what actually led to an observed behaviour. Experiments attempt to identify causal relationships by recreating two or more identical situations in which a causal factor is systematically varied to determine whether or not it will have the prediced outcome. This leads to problems because it can be difficult to ensure that all relevant causal factors have been controlled between the different conditions. Experiments may also have limited 'ecological validity'. This prevents conclusions from being generalised beyond the laboratory into the real world. For example, a study may focus on an analysts ability to use Management Oversight and Risk Trees (MORT) or a similar technique in a silent room without the interruptions that would punctuate their work in an office. In consequence, observational techniques cannot easily be used to provide objective, quantitative comparisons between different reporting systems. They can, however, provide rich insights into the way in which different systems can influence reporting behaviour in complex working domains.

The Ladbroke Grove rail inquiry provides considerable into the potential application of observational and workplace studies to monitor the operation of incident reporting systems. It expressed considerable concern over examples of poor communication and record keeping during the analysis of SPAD incidents. The report argued that 'it is essential that an organisation has a system to record what it has learned, and a process to pass those lessons on to its employees' [194]. Railtrack procedure RT/D/P006 specified that the HQ Production Directorate should monitor and record the implementation of each recommendation. A record could only be closed once the corresponding recommendation had been fully implemented. The Formal Inquiries Process Manager was responsible for following up those recommendations that were directed to Railtrack Headquarters. Although his job description 'clearly envisaged that the progress of recommendations would be tracked after their allocation to individuals, (the Formal Inquiries Process Manager) told the Inquiry that he was given guidance by his managers to the effect that he was not responsible for ensuring that a recommendation was acted upon, but simply that someone had accepted responsibility for it' [194]. The Inquiry concluded that no-one assumed responsibility for monitoring the implementation of recommendations. Cullen also observed that had it not been for the accident and the associated investigation then the shortcomings for tracking incident recommendations might not have been discovered. Faced with this analysis, a new recommendation 'clearing house' was established to collate, prioritise and monitor the implementation of any proposed changes. An important responsibility of the new organisation was to report directly to the Board of the infrastructure company every four weeks.

This analysis indicates the potential application of workplace studies. It emphasises the way in which everyday practice can, over time, depart from published procedures and guidelines. In this example, the responsibilities of the Formal Inquiries Process Manager changed from those documented in the job description and from the intention behind procedure RT/D/P006. This does not imply that such a departure would necessarily have been identified had an observational study been conducted. However, this is precisely the type of working practice that can be observed by these techniques. Hammersely and Atkinson refer to the ways in which the production and use of documents, such as the recommendation reports, are 'socially organised activities' [303]. Ethnographers must, therefore, question the way in which a document is written and distributed. They must also consider what is the purpose and intention behind a document. Ethnographers should also compare the actual use of document against the stated intentions that justify its creation. Differences between observed practice and intended use cannot easily be elicited using monitoring techniques such as focus groups, interviews and questionnaires.

It is important to emphasise that observation techniques can be used to monitor incident reporting systems in situations that extend well beyond the workplace. This is particularly important within the rail industry. Members of the public are often involved in adverse events as well as those who work directly for operating and infrastructure companies. Ethnographic techniques have been widely used to study 'deviant' behaviour. For instance, Popkin et al have recently used this approach to observe patterns of behaviour and control structures within the gangs in many Chicago Public housing developments [683]. Such work provides insights into the relationships between drug use, vandalism, trespass and violence. Many of these activities can have an impact on rail safety. For instance, Smithsimon has conducted a prolonged study of graffiti writers. He argues that ethnographic techniques provide one of the few effective techniques that can be used to gain insights about the behaviours of these individuals and groups. He stresses that 'running from the cops, using the right language, wearing the right clothes: like other ethnographic studies, the right signals and actions, even by an outsider, help gain access to graffiti writers' [746]. Only in this way are 'respondents' willing to provide information about their work and discuss the law-breaking that is a prerequisite for many of their activities. This participation is essential to gain the confidence of individuals who often 'hide behind' the image of a street-wise 'outlaw graffiti artist'. The insights provided by Smithsimon's work can be illustrated by the following except:

> "Hasp and some friends of his offered to show us other, illegal graffiti on the walls lining the adjacent railroad tracks. But he refused to go onto the tracks while a truck belonging to the railroad was parked between the Phactory and the rail yard. John, the photographer, suggested that the railroad employee in the truck probably would not care if we walked down the tracks, but Hasp explained that the Phun Phactory has had repeated problems with the railroad and the transit authority... Before we could learn more about the tension between graffiti proponents and opponents, the truck moved and we traipsed down the tracks, looking at murals. As I spoke with Seac and Ker, they pointed out well-done murals along the walls... 'Get away from there!' yelled an angry voice. Someone on the bridge glared down at us, then dashed away. Hasp and the other writers told us it was someone from the MTA's vandal squad, which focuses on pursuing graffiti artists. We started heading toward the Phactory to get off railroad property. We were about three blocks from the street the Phun Phactory was on, and the row of warehouse walls and razor wire fences along the train tracks meant that if a cop were to get to that street (where the MTA truck had been parked earlier) before we did, we would have been trapped... Meanwhile, I opened my notebook and wrote down the names and descriptions of the writers I had met during the day. 'You writing this down?' asked Hasp. 'What?' I asked. 'What are you writing? This story?' he asked me. 'Oh, no. I'm just writing down everybody's names, and stuff like that'. I flashed a nearly blank page of the notebook toward him, too quickly for him to read much. 'Oh. OK. Cause I thought you were writing this down. Don't write down this', he said..." [746]

Smithsimon's work provides important insights into the behaviour of the graffiti artists who he observed. These insights go beyond the information that can be obtained from incident and accident reports. In particular, it can provide information about potentially dangerous behaviours that are not observed by railroad employees and are, therefore, not reported. Ethnographic studies can also provide insights into the attitudes and shared values that motivate individuals or groups who are involved in trespass or vandalism. For example, Smithsimon argues that 'graffiti represents people's desire to assert their presence in the world through pictures, words, and artistic interpretation' [746]. This conclusion arguably captures the strength and weakness of ethnographic techniques in this domain. It can be difficult to go from the insights that they provide to the recommendations that might avoid future incidents or mitigate the consequences of potential accidents.

Vandalism and trespass are not the only forms of 'deviant' behaviour to be investigated using observation techniques. For example, these approaches have yielded valuable insights into incidents that involve intersections between the road and rail systems. Several studies have shown the difficulty of conducting other forms of investigation into driver behaviour [638, 856]. Individuals will typically take fewer risks and are more likely to obey 'the rules of the road' if they believe that they are being observed. Experimental studies, therefore, seldom yield the violations and extreme behaviours that are witnessed in other contexts. For instance, Burnham's recent study in Alabama observed the behaviour of 862 vehicles as they approached STOP signs at railroad-highway grade crossings [117]. 18% came to a full stop, 50% made a slow rolling stop and 32% did not stop at all. These observations have been interpreted as showing that the majority of drivers do not understand the meaning of the symbols that are used to indicate railroad-highway grade crossings [117]. This is a strong conclusion; an alternative interpretation is that the majority of drivers understand but deliberately choose to ignore the warning signs. Burnham concluded that 'one of the most widely recognised and often overlooked traffic safety axioms is the principle that over use provokes abuse... for a traffic control sign, signal, or pavement marking to be of value it must not be overused' [117]. The difficulty of interpretating observations is a considerable barrier to the practical application of these techniques. Many ethnographers deliberately avoid the 'constructivist theories' that might explain such observed behaviours [303] Unfortunately, these explanations are often essential if we are to be confident in generalising insights from previous failures to predict the likelihood of future incidents and accidents.

It is important to stress that the difficulty of interpreting observed behaviour does not sacrifice the utility of workplace studies and ethnographic techniques. These approaches can often yield important insights even though the complex mechanisms that affect human behaviour are not made explicit. This is best illustrated by the way in which observational approaches can be used to analyse the effectiveness of recommendations that are proposed in the aftermath of adverse events. Again, there are problems with using focus groups or experimental techniques to evaluate these proposals. Expressed opinions may not predict actual behaviour, laboratory conditions may not control all of the factors influencing decision making and performance. Observational techniques provide more direct insights into the 'real world' benefits of potential safety devices. For instance, drivers and pedestrians have been killed and injured by incidents in which they stopped for a first train but then failed to wait for a second train to cross at a junction. The Maryland Mass Transit Administration (MMTA), therefore, tested a 'second train warning' system [523]. This was based around a sign that was illuminated shortly after the first train passed if there was another train approaching. The system was tested at for a ninety day evaluation period at one crossing in Timonium, Maryland. An independent evaluator assessed the performance of the sign by observing driver behaviour and by analysing videotapes. The study concluded that 'risky' driver behaviour decreased by 36% after the installation of the system. Such behaviours can be defined in terms of a specified minimum safe interval between the moment when a vehicle enters the junction and the time at which the second train arrives. A similar study conducted in Los Angeles also used video tape observations on a 'live site' to demonstrate a 14% reduction in risk behaviour. In this instance, 'risky' behaviours were defined to occur when a pedestrian entered the track area six seconds or less before the train entered the crossing [439].

Another major problem complicates the application of observational techniques to monitor the performance of incident reporting systems. In many applications, there are relatively few 'serious' adverse events. In consequence, it is unlikely that an ethnographic or work place study will observe such an incident or accident. These techniques can still provide insights into more frequent, less critical events. However, it is also possible to use some of the analytical techniques that are associated with workplace studies in a post hoc manner. For example, Law has used this approach to demonstrate that 'the character of explanation and cause is relevant in thinking about safety-critical socio-technical systems such as railways' [476]. His analysis of the Ladbroke Grove report and enquiry focuses on a 'rhetoric of spatiality'. Many of the questions and responses during the investigation referred to location, such as 'where does responsibility lie?' or 'Thames Trains could be prosecuted if an incident occurred where driver error was partly to blame'. This use of language reveals that the analysis of failure is understood in terms of distinct 'pigeon-holes' or 'compartments' that are associated with technical, managerial or psychological domains. He identifies other forms of spatial reference. For example, failures can be 'located' within particular subsystems. These views are criticised. Systems responses are compromised by the way in which many social systems are incomplete and unstable. Law argues that in many cases 'the world is simply too fluid and disarticulated' for failures to be located within particular systems. The techniques that Law uses are very similar to those exploited by other sociologists to directly analyse the observations derived from workplace studies. What makes his approach different is that instead of working from his own field notes, his analysis is 'grounded' in the documents produced by an investigation. He is, therefore, sensitive to the context in which such documents are produced. They cannot be analysed at face value but must be seen as publications that are intended to achieve particular objectives.

In passing, it is important to note that Law's ideas have important consequences for analysis of causation in incident and accident reports. His spatial rhetoric can be used to draw conclusions that are broadly similar to many of the other authors in this area who were introduced in Chapter 11. Law criticises the idea of blaming a driver or even the safety culture in an organisation because these are regional interpretations. They place responsibility in a specific pigeon-hole and assumes that blame can be confined within particular boundaries. If the safety culture is at fault then operators can be absolved? However, Law extends his analysis to identify weaknesses in the systemic view of failure. Many of these criticisms have been implicit in the previous chapters of this book, for example in the analysis of some of the findings from the NASA missions in Chapter 10. Law argues that the concept of systemic failure often erodes the boundaries between locations but also often implicitly relies upon the idea that there can be a single focus for particular activities. The proponents of this view, he argues, often talk about 'bringing the system together' or of 'a meeting of minds', For much of the time 'the ordering of the railway is indeed imagined and performed in terms of a system with a more or less strong centre' [476]. This view is, however, flawed. Law argues that the rail system best viewed as a system of dynamic and changing relationships that cannot easily be ordered in such a manner:

> "This is that speed and rapid change together push towards tightly-coupled systems with dense webs of self-sustaining relations. But such systems are best avoided in safety-critical locations. This is because, as we have seen, when things go wrong disruption is rapidly and unpredictably transmitted through the system. Failsafe mechanisms and the tight control of centralised management may work most of the time. But sometimes they will fail. And then they fail there is no play. No slack. Everything falls down. The conclusion is that partially connected, multiply ordered, ambiguous and not very coherent systems are usually more robust. And the corollary is that if we find that we are proposing technologies that demand tight systems then we need to stop and think. This the ultimate lesson of the Ladbroke Grove tragedy. It is that we have unwisely created a world which demands coherent systems" [476]

It is interesting that sociological approaches should at the same time yield immensely detailed observations of group and individual behaviour as well as such high-level insights into safety-critical organisations. Unfortunately the broad range of these approaches cannot overcome some of the problems that arise when attempting to use the resulting insights to improve safety. Workplace studies and sociological analyses seldom yield direct recommendations. In many ways, this is the point behind the techniques. The insights they provide inform decision making but do not automatically help to shape or focus those decisions. In contrast, statistical techniques can be tailored more directly to support particular hypotheses about reporting behaviour and the operation of reporting systems.

15.4.5 Statistical Analysis

Previous paragraphs have introduced different forms of statistical analysis that support the monitoring of incident reporting schemes. These include simple incident and reporting frequencies as well as threshold models, such as UK Railway Safety's Precursor Indicator Model, and more advanced techniques, including least squares regression used for trend analysis [697]. Several specialist textbooks provide an introduction to the particular mathematical approaches that support these techniques []. In contrast, the remainder of this section focuses on the managerial and organisational issues that must be considered when using statistical methods to support the monitoring of incident reporting systems. For instance, it is important to consider whether particular numerical values can yield 'valid' insights into the performance of the underlying systems. It is for this reason that the Transportation Safety Board of Canada do not preset accident totals for particular railways [787]. They argue that 'the track, rolling stock and personnel in an occurrence may all belong to different companies; also an occurrence may have several contributing factors'. Presenting data about one of these factors might be 'misleading' and there is a danger that misinterpretation of the data could have an unfair affect on a company's 'competitive position'. As we have seen, other organisations reject this argument and instead rely upon normalisation techniques to help make valid comparisons between different organisations. For instance, the independent review of Australian rail safety argued that without 'measurable, appropriately normalised data' it is impossible to determine whether the industry is becoming safer; whether passenger safety is improving or not and whether there are significant trends in freight and passenger train incidents and accidents [55]. Table 15.13 illustrates this point. It documents the number of fatalities associated with different modes of transport and is taken from the the Australian rail report cited above. It is difficult to make direct comparisons between these statistics because the table does not record the risk exposure associated with each mode. For example, the relatively high number of fatalities associated with road travel can be offset against the disproportionately large number of journeys or trip distances that are made each year using this form of transport. Similarly, the low number of deaths from maritime transportation cannot be correctly interpreted without information about the numbers of people involved in this industry.

	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	% Fall
Rail	60	68	76	48	61	52	39	48	34	39	-35%
Road	3197	2935	2601	2324	2172	2048	2044	2126	2031	1876	-41%
Water	69	60	92	89	77	73	62	60	57	48	-30%
Air	57	97	74	65	79	87	50	65	71	49	-14%
Other	6	3	6	2	1	0	1	2	1	2	
Total	3389	3163	2849	2528	2390	2260	2196	2301	2194	2014	-41%

Table 15.13: Transport Accident Deaths, Australia 1988-1997 by Year and Mode [55]

Unfortunately, normalisation does not offer a panacea for the problems of interpreting incident statistics. As we have seen, it can be difficult to agree upon appropriate criteria. For example, the information in Table 15.13 might be normalised to present the ratio of deaths to passenger miles. This would clearly be appropriate within mass transportation modes such as the road and rail systems. It is less certain that this normalisation would yield meaningful data for the maritime industries. Relatively small numbers of passengers are carried; fatalities are often associated with shore personnel servicing shipping. It might, therefore, be more appropriate to normalise according to the tonnage carried within each industry. This would raise further problems in the analysis of rail and road data unless a distinction was made between fatalities involving freight and passenger services.

The choice of normalisation criteria can have an important impact on the calculation of trend statistics. Different industries respond in different ways to changes in the underlying economic cycle. For instance, adverse conditions in the world economy during 2001 have arguably had a more profound impact on air transport than they have upon road transportation. Normalisation factors that ignore the impact of this down-turn upon passenger traffic might, therefore, exaggerate any associated drop in the incident rate. Similarly, as with many forms of statistical analysis, trend identification can be profoundly affected by the base date that is used in any comparison. It can be difficult to interpret the significance of the percentage reductions in Table 15.13 without understanding the reasons for selecting 1988 as the starting point.

The problems of normalisation and of trend analysis are general in the sense that they affect the monitoring of many different systems. A number of further problems specifically affect the monitoring of incident and accident statistics. Many of these stem from the causal asymmetries that have been described in previous chapters. Statistical returns often depend upon subjective analysis to determine the causal factors behind past events. Changes in aggregate data can reflect changes in interpretation rather than new forms of failure within an industry. In Table 15.13, problems arise when compiling statistics for incidents involving more than one mode of transport. As we have seen, many rail incidents involve pedestrians and road vehicles. It is difficult to interpret the statistics provided in this summary without some explanation of how such accidents would be encoded. Such fatalities might be associated with the rail system or with the road system or both.

There are a number of specific problems that complicate the compilation of incident and accident statistics. It is often difficult to know whether or not particular effects can be attributed to an adverse event. For example, the data is Table 15.13 cannot easily be interpreted without additional information about the definition of a transport-related fatality. This is important because an individual may die several days after an incident has occurred. In extreme cases, they may receive injuries that contribute to their death many months or even years after the adverse event. This attribution problem is exacerbated for occupational illnesses where individuals may also be exposed to other contributory factors within their wider environment. In the UK, this reasoning led to the Court of Appeal 'Fairchild' decision (December 11, 2001, Lord Justices Brooke, Latham and Kay). This focussed on cases brought by victims of the mesothelioma lung disease. Mesothelioma is linked to exposure to asbestos products. The defendants were all employers or operators of premises where asbestos was being used or cleared. The Appeal Court Justices refused damages to the claimants because mesothelioma 'is a single indivisible disease and a claimant cannot establish on the balance of probabilities when it was he inhaled the asbestos fibre, or fibres, which caused a mesothelial cell in his pleura to become malignant'. The impact of this ruling has been profound. It has subsequently been argued that since it is impossible to identify the individual fibre that causes the mesothelioma then any exposure to asbestos should be regarded as a possible cause. From this it follows that the level of liability should reflect the degree of exposure to any potential cause of the disease. It remains to be seen whether this line of argument will stand against the Fairchild decision.

Incident and accident statistics can only be interpreted correctly if analysts understand the criteria that guided the collation of source data. Using different definitions for reportable incidents can lead to very different statistics being presented. This can be illustrated by the way in which the Transportation Safety Board analysis the performance of its rail network using both Canadian criteria and the criteria proposed by the US Federal Railroad Administration requirements. The Canadian criteria since 1992 consider that all main-track and non-main-track accidents are reportable as long as the damage to rolling stock renders it unsafe. The Federal Railroad Administration requires a minimum dollar damage threshold of $6,300 US for all reportable accidents. This policy of using dual criteria in the collation of accident statistics enables accurate comparisons to be drawn between these two different approaches. For instance, the data compiled by the Canadian Pacific Railway for the January - August period of 1994-1996 show very different trends in main track derailments depending on the criteria used.

> "... when TSB reporting requirements were used, CPR's main line derailments were 30% higher during the January - August period of 1996 than during the corresponding period in 1994 and 1995. However, when FRA reporting criteria were used, the number of main line derailments remained unchanged throughout the period being examined. During this three year period, an average of 75% of occurrences that met TSB but not FRA accident reporting guidelines involved derailments of only one car." [783]

The maintenance of different statistics to reflect different reporting criteria is a general problem. For example, it affects many agencies and commercial organisations that implement their own local criteria but must then follow different agreed criteria when reporting to higher organisations. For instance, most European Air Traffic Management organisations must pre-process their incident statistics before submitting them to EUROCONTROL [423]. One consequence of this is that analysts must always check which criteria are being applied when interpreting meta-level trend information. There are additional complexities. For instance, reporting systems often revise their own criteria. This creates problems when analysts attempt to draw comparisons between more recent statistics and those gathered under previous reporting criteria. For instance, the Canadian Transportation Safety Board revised its guidelines in August 1992. Before this time, derailments and collisions were only reportable if casualties or dangerous goods were involved or for main-track accidents if there was property damage in excess of a monetary threshold. As we have seen, since 1992 all main-track and non-main-track accidents are reportable as long as the damage to rolling stock renders it 'unsafe'. After 1992, all crossing accidents are reportable. Prior to that year accidents at farm and private crossings were reported only if they involved a casualty/dangerous goods/derailment resulting in property damage in excess of a monetary threshold [787]. It is difficult to underestimate the consequences of such changes on the compilation of incident statistics. Some occurrence categories previously regarded as incidents were now regarded as accidents. Other types of occurrence were no longer reportable. The changes also made it difficult to calculate trends, data reported under the new definitions could not be directly compared to historical data that was gathered under the previous criteria. Where possible the Safety Board revised previous data in an attempt to adapt it to the new criteria. They did, however, emphasise that 'caution is required' when comparing statistics before and after the reporting requirement change and that 'the interpretation of the results from recent years has been clouded because of the change in TSB reporting requirements'.

The problems of gathering, of normalising and of interpreting the statistical information used to monitor reporting systems has led some safety-related organisations to develop extensive criteria to guide many different aspects of data analysis. For instance, the UK Health and Safety Executive have been involved in an initiative to 'revitalise' the use of Health and Safety targets to promote national performance. Part of this work has involved the development of a 'note' to ensure the validity of the statistics that will be used to measure progress towards these revised objectives [337]. This note lays out a number of general principles:

1. "Progress measurement will involve more than one data source and some adjustment or integration of data from the different sources will be necessary; as a rule this will only be appropriate at the global level.

2. Percentage changes over time are what matter for monitoring progress against the targets, so efforts should be focused on measuring change; estimates of absolute levels may vary as information sources evolve.

3. In assessing trends and progress over the strategy period, statistical modelling techniques will be used to limit the impact of sampling variability in the figures for individual years.

4. To support the outcome data on injuries and ill health, supplementary approaches should be explored, for example collecting data on economic, social and cultural factors.

5. The data and methods used for progress measurement will be National Statistics, so the methods will be subject to independent quality review and stakeholder consultation.

6. A report on progress will be prepared each autumn, comparing the latest data with those for the base year (1999/2000). For at least the mid- (2004/5) and end-point (2009/10) of the strategies, this report will incorporate external peer review." [338]

The note also includes specific sections describing techniques for describing the injuries target. For instance, the reporting rate for less serious injuries will be used to calculate an adjustment for the under-reporting of major injuries. It also lays out criteria for the statistical analysis of work-related ill health. Existing data sources will be refined, for example to account for 'the effects of raised awareness'. The note also promises to identify new sources of statistical data, including workplace-based surveys. Diseases with long latency periods between exposure and health outcome will be included but will be handled separately from other illnesses for the reasons described in previous paragraphs.

The note builds on International Labour Organisation recommendations by arguing that data from different sources should be integrated 'to produce an overall judgement about progress'. The note also outlines some of the high level problems that arise from this approach. The integration of data can often involve labour intensive adjustments, for instance where they may be subtle differences in the periods over which aggregate data has been compiled. Also 'where one or more of the sources involve sampling, the reliability of the resultant estimates will reduce as the level of disaggregation increases' [338]. There are further technical problems. For example, statistical measures are subject to random error. It is, therefore, usual to indicate a central estimate for any measure together with upper and lower confidence levels around this estimate. It, therefore, follows that the lower confidence limit must equals or exceed the target value before analysts can argue that the target has been achieved. Similarly, a target can be shown not to have been met only if the upper confidence limit falls short of that target. If the target value falls between the upper and lower confidence limits, no definitive statistical judgement can be made. These importance of confidence levels is often underestimated or ignored by analysts who interpret incident statistics. It can be argued that this book also falls into this trap by postponing any discussion of these issues until relatively late in our exposition. The HSE emphasise that these issues must be considered precisely because even our best measurements are subject to 'uncertainty', to 'under reporting' and to 'sampling error'. For example, a 1999 labour force surveys reported that there were 380,000 reportable injuries to workers. Of these, 343,000 were to employees. Employers, however, only completed 161,000 injury reports. Self-employed people made 1,599 non-fatal injury reports in 1998/99. This compared with 35,000 injuries estimated by the same survey. This suggests a reporting level of less than 5% for the self-employed. Similarly, the HSE report that the margin of uncertainty on disease estimates drawn from self-reported surveys is up to 30% for stress/depression/anxiety, upper limb disorders and back disorders. The margin of error is assumed to be lower for data collected using Sentinel systems, such as the FDA's trials mentioned in Chapter 14. However, the HSE report the lack of any agreed methodology for measuring sampling error in these systems. It may not, therefore, be possible to use statistics to determine whether or not an incident reporting system has actually met a particular target!

The statistical note, mentioned above, argues that the relatively large statistical errors associated with accident and incident reporting represent an 'unsatisfactory situation' and urge analysts to 'reduce the statistical uncertainty to as low a level as possible' [338]. Partial solutions include the use of statistical modelling across several years. Data can be taken from successive years rather simply comparing the base and final years of the sample. Overlaps between samples for successive periods can also be used [697]. The HSE argue that 'the precise statistical models will depend on the series that emerge'. A uniform decline in incidence rates would justify the use of simple linear regressions. More complex methods may, however, be needed to explain complex trends. This would be the case if an initial fall in incidents rates was not sustained. The HSE note that a recent analysis of German accident rates suggested that trends were best modelled as an exponential rather than a linear decline. They emphasise that these decisions must be 'data driven'. They must not be influenced by whether or not they provide a favourable answer [338].

Even if statistical studies obey well intentioned guidelines, such as those cited above, there can be problems in the monitoring of reporting systems. In particular, problems can arise from the degree of sophistication implied by those guidelines. The end recipients of the statistics may fail to understand or interpret the information that they are being provided with. This point can be illustrated by recent problems in the presentation of UK SPAD statistics.

> "HSE statisticians have advised that the method of presenting SPAD information should be revised to avoid potentially misleading interpretation of the standard analysis. Previously, ... the standard presentations each month have been represented using the ratio of that month's SPAD count with the average of the corresponding months of the six preceding years. These data were then plotted out month by month, together with a 'trend' line fitted to these points (technically a linear least-squares squares regression line). The main visual message of this representation is the trend line, which could lead to readers naturally assuming that this represents the trend in SPAD numbers (i.e. if the slope is up, SPADs are increasing, and vice versa). However this natural assumption is wrong. The slope of the trendline indicates whether the change in SPAD numbers (change being measured over the past six years) is getting bigger or smaller month by month, regardless of whether the change itself is upwards or downwards. A flat trend indicates a steady increase or decrease: it does not discriminate between the two. The

> previous presentation thus shifts attention from the issue of primary interest (are SPADs increasing or decreasing?), to a secondary issue (is the rate of change in SPAD numbers increasing or decreasing?)." [354]

In other words, the statistical techniques were appropriate for the data being analysed. However, the graphical presentation of those statistics was easily misinterpreted. Fortunately, there are a number of texts that discuss appropriate presentation formats for such information. For instance, Tufte's books warn analysts about a range of common biases that affect our interpretation of the graphical presentation of statistical information [789]. Rather than repeat this material, the following pages focus on a range of computer-based visualisation techniques that have been developed to support the monitoring of incidents and incident reporting systems.

15.4.6 Electronic Visualisation

The previous section has argued that many safety managers and regulators have difficulty in interpreting the statistics that are collated to support the monitoring of incident and accident reporting systems. Engineers and scientists 'need an alternative to numbers' when analysing such complex data sets [526]. Graphical visualisations can be used to address this problem. However, as the previous quotation illustrates, there are also associated problems when people fail to correctly interpret those graphical representations. Previous arguments can be illustrated by UK SPAD statistics. For example, UK SPADs are categorised according to a severity classification scheme. The following definitions introduce the term 'signal overlap'. This is the distance specifically provided after signals as a safety margin to cater for misjudgement or problems with the train braking systems. Italics are also used to represent changes from previous definitions. The associated HSE reports do not describe these changes in detail. It must be assumed that the HSE have updated the historic data to reflect the new definitions. This further emphasises the importance of providing sufficient information about the criteria used when compiling statistics so that readers can correctly interpret the impact of such changes in the definitions of particular categories:

- Category 0: Not entered
- Category 1: Overrun 0 to 25 yards, *overrun not exceeding overlap*, and no damage, injuries or deaths.
- Category 2: Overrun 26 to 200 yards, *overrun not exceeding overlap*, and no damage, injuries or deaths.
- Category 3: *Overrun greater than overlap* plus all overruns greater than 200 yards and no damage, injuries or deaths.
- Category 4: Track damage only with no casualties.
- Category 5: Derailment with no collision and no casualties.
- Category 6: Collision (with or without derailment) and no casualties.
- Category 7: Injuries to staff or passengers with no fatalities.
- Category 8: Fatalities to staff or passengers.

Table 15.14 presents an extract from the associated UK SPAD statistics. Although it is relatively easy to extract salient features from this data, it is important to remember that it only represents a very limited snapshot of previous SPAD incidents. The introduction of additional information, such as the operating companies involved in each SPAD, can add considerable complexity to the interpretation of these statistics. The visualisation of this additional detail will be addressed in subsequent pages. For now, it is sufficient to observe that the statistical information in Table 15.14 can be visualised in a number of different ways.

	94/95	95/96	96/97	97/98	98/99	99/00	00/01
0	0	0	0	0	0	0	0
1	43	38	39	35	35	38	33
2	28	33	26	28	29	31	24
3	20	24	28	29	28	27	37
4	7	4	5	4	5	2	2
5	1	1	2	2	2	2	4
6	0	0	0	0	0	0	1
7	1	1	0	1	1	1	0
8	0	0	0	0	0	0	0

Table 15.14: Percentage of Total SPAD Incidents by Severity by Year [353]

Figure 15.1 illustrated a 'conventional' visualisation for this data, based on several similar graphs in the HSE SPAD returns. As can be seen, the percentage of incidents at particular levels of severity are mapped onto the Y-axis. The X-axis is used to denote the data set from each year's returns. Lines can be plotted between individual data points to illustrate potential trends between SPAD incidents at particular levels of severity. It is important that safety managers and the administrators of incident reporting systems carefully consider the strengths and weaknesses of such visual representations. On the one hand, Figure 15.1 provides a relatively clear overview of the data in Table 15.14. On the other hand, it fails to capture the confidence intervals that was emphasised as a important part of any statistical analysis in the previous section. This can be done by introducing bars above and below the central point for each data value in Figure 15.1. There are further problems. A far greater percentage of incidents occur at severity levels 1,2 and 3 than 0, 4, 5, 6, 7 and 8. Similarly, the relatively low proportion of incidents at severity levels 0 and 4to 8 also creates considerable overlap between the 'trend lines' in Figure 15.1. It is difficult to distinguish between small differences in the proportion of incidents at these severity levels over the time period concerned. Such problems might be addressed by using two different scales on the Y-axis. A relatively large interval might be used for the large differences between values for severity levels 1,2 and 3. A more fine-grained scale might be introduced for the lower proportion of incidents at severity levels 0 and 4 to 8. It is important to emphasise that these caveats are typical of the problems that complicate the visualisation of incident reporting statistics in many different industries. The severity distribution reported by the HSE is similar to those identified by EUROCONTROL in air traffic management [423] and by maritime reporting agencies [366].

A number of further problems affect the visualisation in Figure 15.1. In particular, readers may fail to notice that the values in this graph represent percentages rather than absolute frequencies for SPAD incidents. This potential pitfall is addressed by the alternative visualisation in Figure 15.2. In this representation, the data is cumulatively mapped onto a percentage scale so that users can see the way in which different levels of severity contribute to that total proportion of SPAD incidents. This visualisation arguably 'exposes' the relatively large percentage of level 8 events in 1994-1995 more clearly than the previous representations. As with Figure 15.1, number of further criticisms can be made. It is difficult to identify the precise percentages for each severity level, especially for the categories 0 and 4 to 8. This could be addressed by introducing labels to provide the numeric values in table 15.14. These might also provide absolute numbers of SPADs in each category to avoid the confusion between proportions and frequencies, mentioned above.

Many further visualisations can also be used to represent the data in Table 15.14. For example, Figure 15.3 uses a form of 'radar' presentation. A separate axis is drawn for each of the data sets being considered. Each axis begins at a common origin and ends at a point such that there is a uniform distance between consecutive axes on the circumference of a circle. Lines can be drawn between the proportion of incidents at a particular severity level for each data set. If the proportion remained the same then one would expect to form a regular shape. In this case, with seven data sets we would anticipate regular heptagons. This is illustrated by SPADs at severity levels 1 and 2 whereas

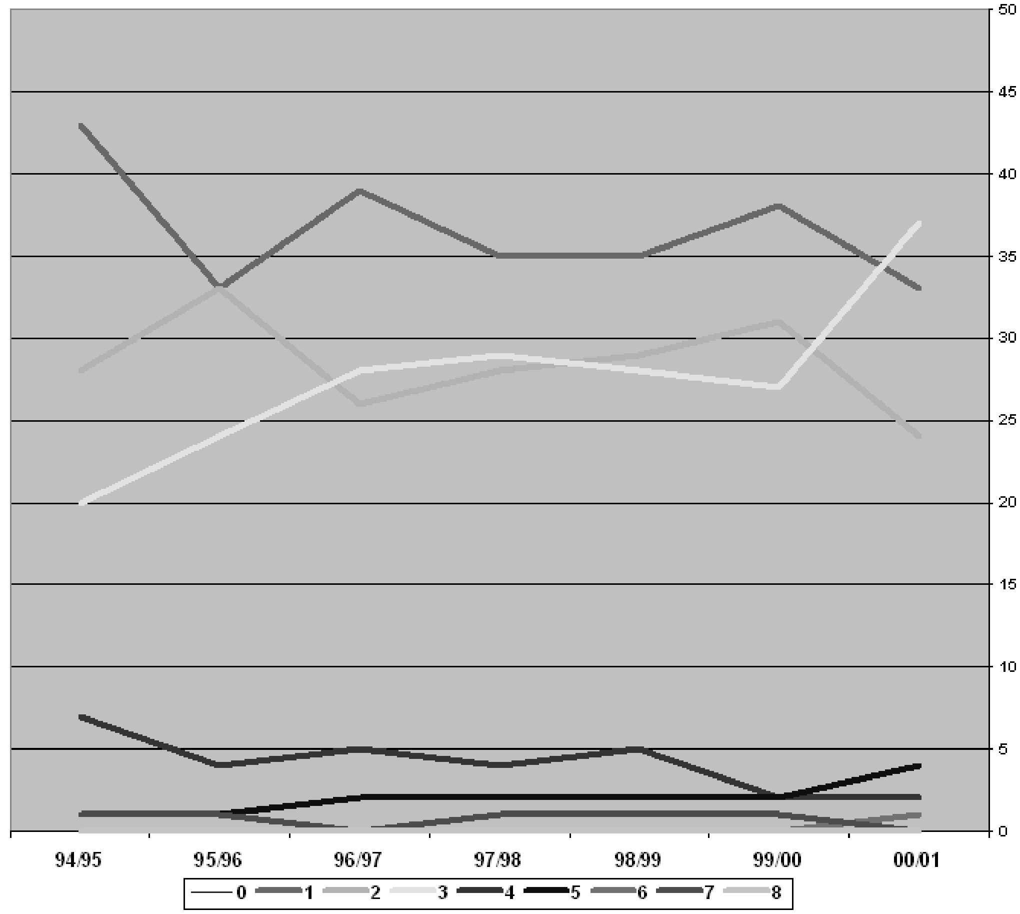

Figure 15.1: Static Conventional Visualisation of SPAD Severity by Year

a clear distortion can be observed at level 3 in 2000-2001. Again, however, the use of a uniform scale along each of the seven axes creates problems in distinguishing the relative proportion of incidents in categories 0 and 4 to 8. The key point in introducing each of these alternate representations is that their strengths and weaknesses must be matched against the particular requirements of their intended users. Figure 15.2 arguably provides a more accurate impression of the relative percentages at particular severity levels. It would, therefore, be appropriate for presenting this data to analysts who were not frequently required to monitor these statistics. In contrast, Figure 15.3 and Figure 15.1 might be used in circumstances where readers might be expected to recognise that these images did not record incident frequencies at particular severity levels .

Most of the visualisations that support the monitoring of incident reporting systems are generated by teams of statisticians who work from incident databases. They publish summary documents that are distributed to more senior management at regular intervals. This approach can create a number of limitations. For example, there may be a significant delay between the collation of incident data and the publication of the graphs and forms that are used to monitor those statistics. This prevents regulators and managers from taking prompt action in response to sudden changes. Electronic visualisation systems avoid these problems by linking graphical representations to on-line incident information systems. Users can automatically update the values in a visualisation to reflect the most

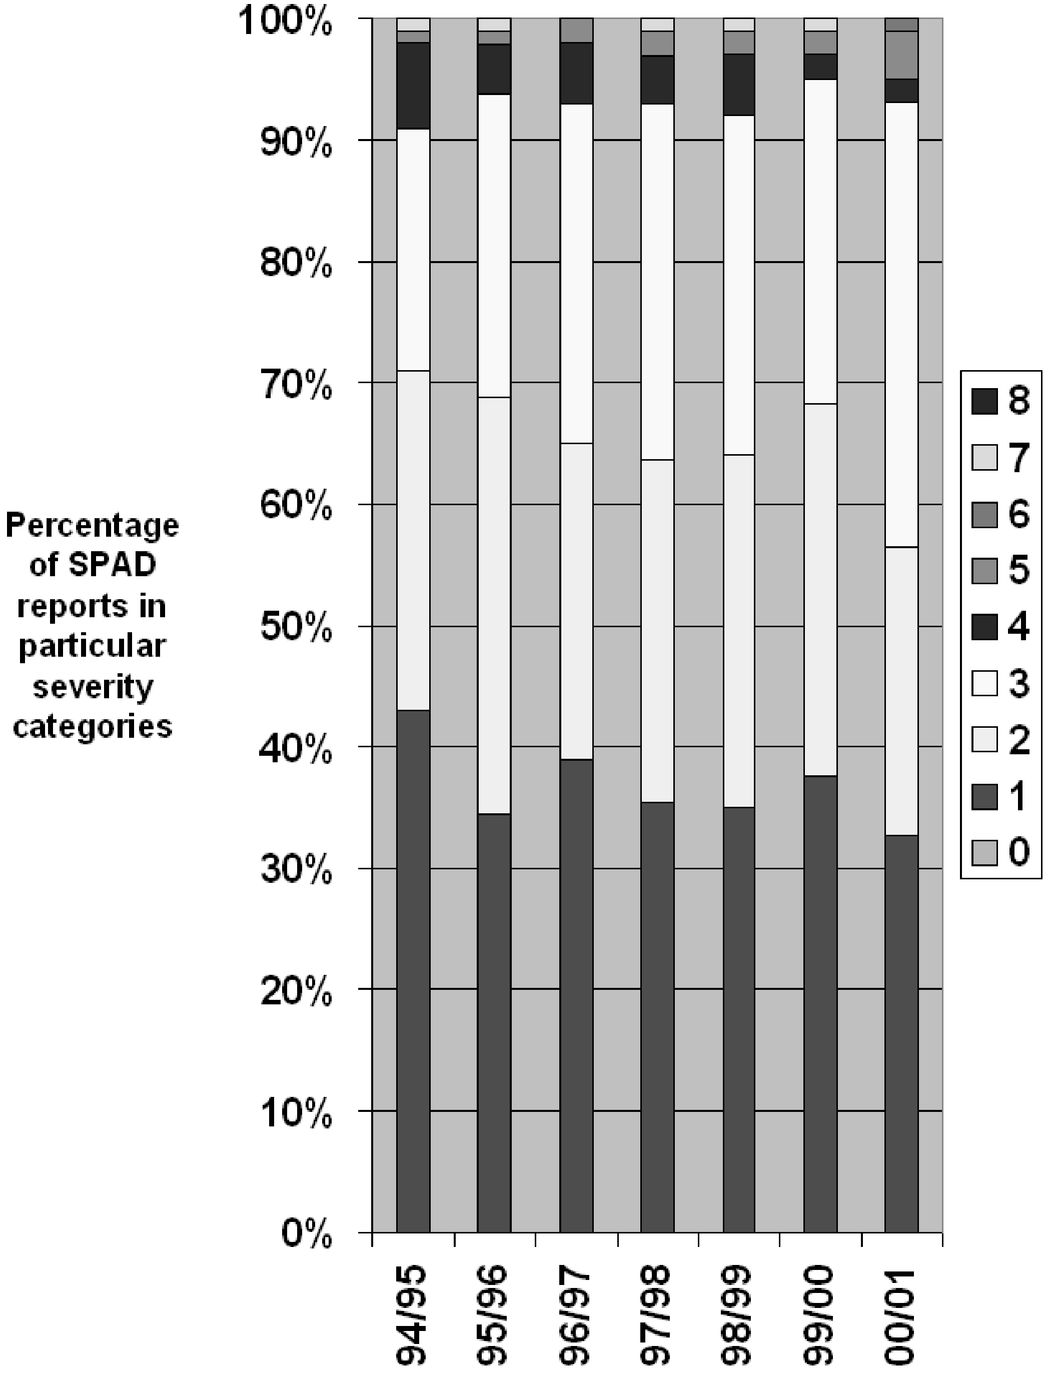

Figure 15.2: Static 'Column' Visualisation of SPAD Severity by Year

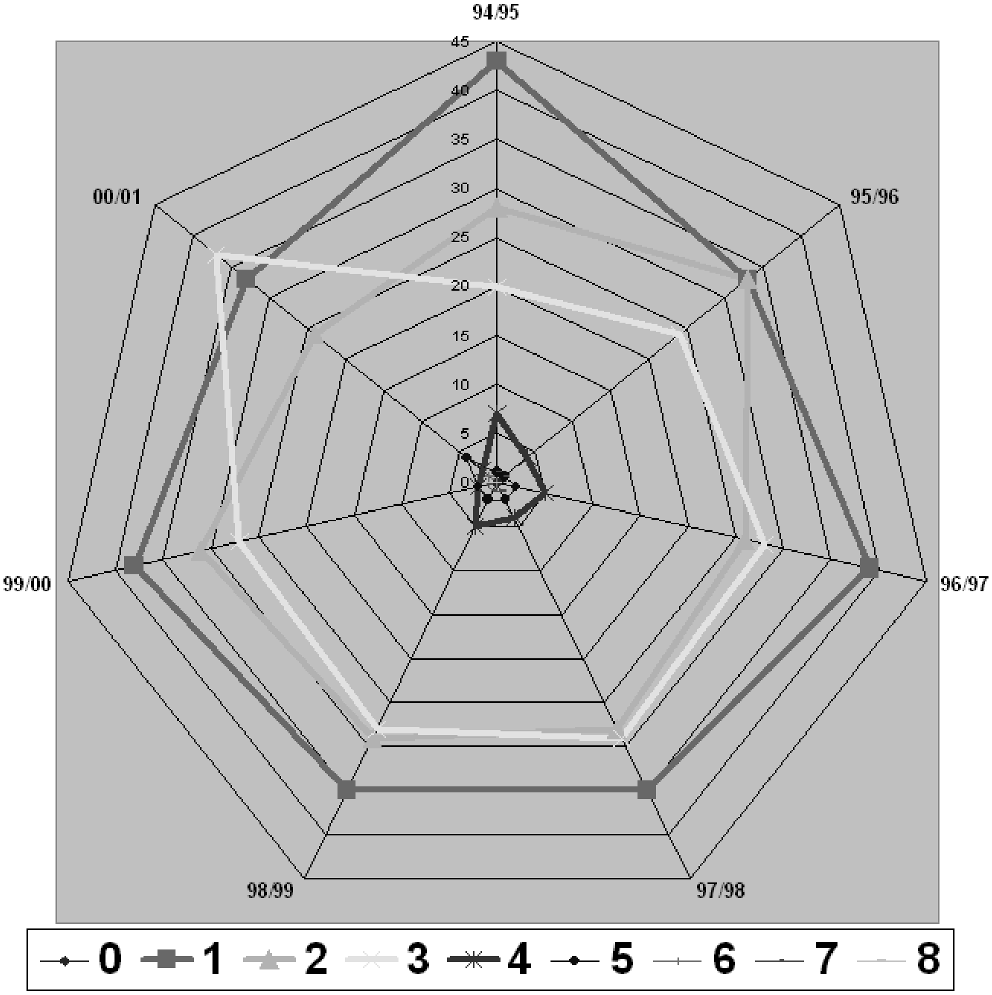

Figure 15.3: Static 'Radar' Visualisation of SPAD Severity by Year

recent changes to a database. There are further speed benefits . The use of wide area networks can rapidly disseminate monitoring information across many different contributory organisations and stake holders. For example, UK Railway Safety's Safety Management Information System (SMIS) provides on-line updates about the 250 safety events that are reported across the network each day. Over 50 organisations and 300 registered users can access updates about these incidents as they are entered into the system.

Electronic visualisation tools also offer a number of further benefits. Safety managers often find it difficult to produce new visualisations that might better help them explore alternative hypotheses about the causes and consequences of adverse events. If standard graphs and bar charts do not satisfy their requirements then managers must negotiate with their statisticians to request changes in the standard presentation formats. In contrast, many reporting systems now provide managers with direct access to incident data through the relational database technology described in Chapter 14. These systems help safety managers to pose interactive requests for data using the Structured Query Language (SQL) and its derivatives. Alternatively, user interface facilities can be provided to simplify query composition by supporting a more limited range of retrieval tasks. In either case, it is possible to use these information requests to drive graphical visualisations of statistics that characterise the information within the relational database. For example, Railway Safety's SMIS integrates with another PC-based reporting package (Crystal Reports). This enables users to customise the extraction of information from the safety management system to drive a large number of alternate visualisations.

As we have seen, however, a number of problems limit the utility of relational systems. It can be difficult for many safety professionals to master the syntax and semantics of SQL queries. For example, the following command for extracts SPAD data from an SQL database described by Speirs and Johnson [752]. It selects incidents where the overshoot is between 250 and 1,000 yards and the severity level was between 4 and 8 occurring between 18:00 and 24:00. The syntax has been simplified from that used in the full implementation of the reporting system!

```
select * from Incidents as i
      where Overshoot between 250 AND 1000
      AND ProvisionalSeverity between 4 and 8
      AND Time between 180000 AND 240000
```

The problems that managers have in interpreting relational commands has profound implications; 'inaccurate' queries can trigger recommendations that might not otherwise have been made. The use of pre-canned queries developed by support staff can be equally problematic. Previous studies have shown that safety managers often misunderstand the information requests that others develop for them [748]. Further criticisms can be raised because many existing electronic visualisations are simple extensions of paper-based representations. They are dominated by bar charts, scatter plots and conventional graphs. This is a missed-opportunity given the more advanced visualisation facilities that have been developed in other application domains. A further problem is that standard bar charts and graphs cannot easily be used to represent the complex 'multi-dimensional' data that is gathered about adverse events. For instance, Table 15.15 provides a brief excerpt from the SPAD data that is routinely disseminated by the HSE. It is unclear how any single visualisation might be used to capture this heterogeneous information. If such a system could be developed then users might be enabled to explore different attributes of the data. For example, it might be possible to explore whether a particular zone suffered from incidents at a particular level of severity at a similar time of the day or week. Alternatively, it might be used to determine whether the consequences of SPADs in some zones were effectively mitigated by defensive driving, leading to lower severity incidents with limited overshoots.

Previous paragraphs have criticised many incident reporting systems because they have failed to exploit the benefits offered by recent advances in information visualisation. A vast range of computer-based visualisations have been developed to help users extract statistical information from medical and other forms of scientific data [526]. They have also been used to visualise document collections and navigate the many thousands of potential results from web-based search [706]. This work has had little impact on the current generation of relational databases that support most incident reporting systems [753]. Speirs [752] has, therefore, begun to transfer more recent visualisation technology to support the presentation of UK SPAD data based on the subset of SMIS, presented in Table 15.15. The intention is not that these visualisations would replace other forms of numerical data analysis. In contrast, the aim is to provide a decision support tool that will enable safety managers and regulators to perform more interactive forms of search. Such visualisations are intended to help users 'discover' new properties of incident data that will enable them to monitor both changes in the underlying failures and the performance of the reporting system itself.

The design of the prototype visualisation was based around requirements derived from discussions with staff involved in the original, UK CIRAS confidential incident reporting system. These discussions emphasised the importance of juxtaposition within any visualisation for rail incident data. For example, it is important to compare the incidence of multiple and single SPADs within the same region. This information can be used to identify existing 'hot spots' where repeated SPADs have occurred in the past. This information can then be used to anticipate future problems where single incidents might, over time, become multiple SPADs. It is important to emphasise that one of the aims in implementing the SPAD visualisation tool was to identify the requirements for future interfaces to incident data. Before building such a prototype, it was difficult to provide railway

Date (2000)	Time	Signal No.	Location	Train Operating Company	Zone	Dist. Passed (Yards)	No. of SPADs at signal	SPAD severity	No. of SPADs by driver	HMRI Action Level
1/11	0339	B248	Narroways Junct.	EWS	GWZ	2377	8	7	2	3
1/11	1910	LD32	Liskeard	Wales & West	GWZ	12	2	3	1	2
1/11	1310	TT83	Pye Bridge	EWS	MZ	30	1	2	2	2
1/11	0430	T626	Horsham Road Xing	AMEC	SZ	200	1	3	1	3
2/11	1058	BW6	Bottesford West	Central Trains Ltd	MZ	200	1	2	1	3
4/11	1015	RETB	Halesworth	Anglia	EAZ	50	1	2	1	1
5/11	1705	St And X	Bristol Temple Meads-P5	Wales & West	GWZ	85	3	3	1	2
5/11	0030	LR534	Hathern	EWS	MZ	200	1	2	2	2
7/11	1035	HP20	Harringay Park	Silverlink	EAZ	1056	3	3	3	3
7/11	1240	E118	Taunton	Wales & West	GWZ	25	3	1	2	2
8/11	0620	R838	Newbury	Thames Trains	GWZ	3	3	1	3	3

Table 15.15: Subset of SPAD Incidents for November 2000 [353]

staff with an accurate impression of what was, and what was not possible, given current technology. The prototype has since been evaluated with support from UK Railway Safety and this process has helped to provide more detailed requirements for future versions of the visualisation. Some of the more surprising results from this consultation process are discussed in later paragraphs.

Figure 15.4 provides a screen shot from an initial version of the visualisation tool. The top left of the screen presents a map of the UK, each dot on the map represents the location of a SPAD. Colour coding can be used to distinguish multiple SPADS. The user can select each of these icons to obtain a range of more detailed information about the signal location. Figure 15.5 shows how this can include a time-line of previous events relating to the placement of the signal and also any previous SPADs. It can also include photographic information as well as plans and 3D models of the signal location.

The bottom of the screen in Figure 15.4 provides access to the more detailed information that the system holds about each SPAD. This corresponds to an extended form of the information presented in Table 15.15. By default, every SPAD incident is represented by a single dot and by a row in the panel at the bottom of the screen in Figure 15.4. However, the exact number of dots and rows will change in response to user selections. These are formed using the panel on the top right of the display. Figure 15.6 provides a more detailed image of these input widgets.

The visualisation tool uses a technique known as 'dynamic querying' [6] to avoid the problems associated with forming and interpreting the results of SQL statements. This technique enables users to extract information from a data set by directly manipulating common interface widgets such as sliders, lists and radio buttons. In Figure 15.6, the user can select either end of the Overshoot, Severity Category, Time of Day and Date sliders. For example, if they select the left hand icon on the

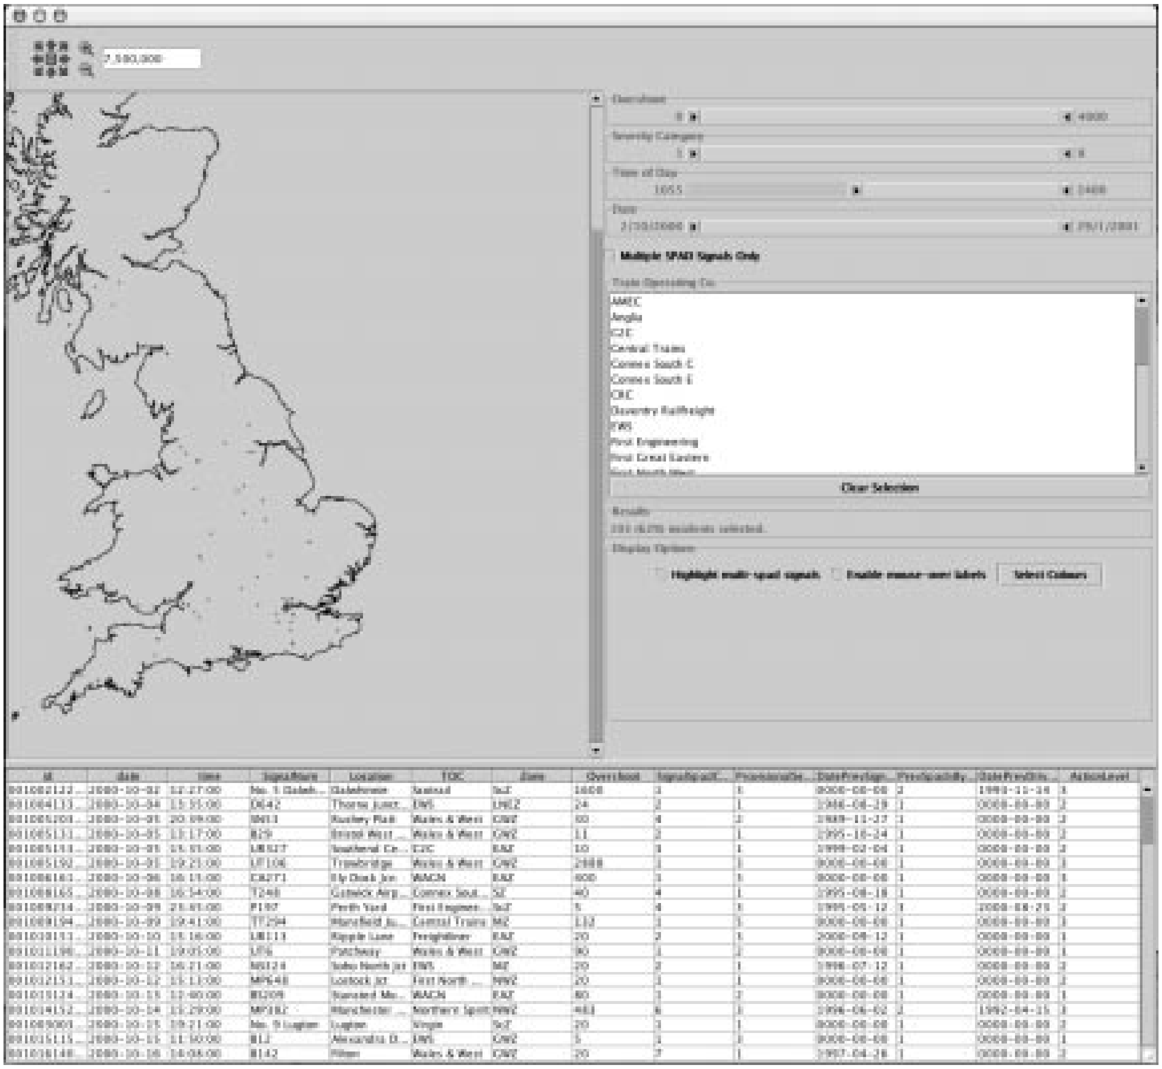

Overshoot: 0 – 4000

Severity Category: 1 – 8

Time of Day: 1055 – 2400

Date: 2/10/2000 – 29/1/2001

□ Multiple SPAD Signals Only

Train Operating Co.

AMEC
Anglia
C2C
Central Trains
Connex South C
Connex South E
CRC
Daventry Railfreight
EWS
First Engineering
First Great Eastern
First North West

Results: 231 (6.0%) incidents selected.

Display Options: □ Highlight multi-spad signals □ Enable mouse-over labels

id	date	time	SignalNum	Location	TOC	Zone	Overshoot	SignalSpadC...	ProvisionalSe...	DatePrevSign...	PrevSpadsBy...	DatePrevDris...	ActionLevel
001002122...	2000-10-02	12:27:00	No. 1 Galash...	Galashmie	Seatrail	ScZ	1608	1	1	0000-00-00	2	1993-11-16	3
001004133...	2000-10-04	13:55:00	D642	Thorne Junct...	EWS	LNEZ	24	2	1	1986-08-29	1	0000-00-00	2
001005203...	2000-10-05	20:09:00	BN13	Bushey Park	Wales & West	GWZ	00	4	2	1989-11-27	1	0000-00-00	2
001005131...	2000-10-05	13:17:00	B29	Bristol West ...	Wales & West	GWZ	11	2	1	1995-10-24	1	0000-00-00	2
001005153...	2000-10-05	15:55:00	LB527	Southend Ce...	C2C	EAZ	10	3	1	1999-02-04	1	0000-00-00	2
001005192...	2000-10-05	19:21:00	UT106	Trowbridge	Wales & West	GWZ	2888	1	1	0000-00-00	1	0000-00-00	3
001008161...	2000-10-08	16:15:00	CA271	Ely Dock Jcn	WACN	EAZ	600	1	1	0000-00-00	1	0000-00-00	3
001008165...	2000-10-08	16:54:00	T248	Gatwick Airp...	Connex Sout...	SZ	40	4	1	1995-08-18	1	0000-00-00	2
001009214...	2000-10-09	21:45:00	P197	Perth Yard	First Enginee...	ScZ	5	4	1	1995-05-12	1	2000-08-21	2
001009194...	2000-10-09	19:41:00	TT294	Mansfield Ju...	Central Trains	MZ	132	1	1	0000-00-00	1	0000-00-00	3
001010151...	2000-10-10	15:16:00	LB113	Ripple Lane	Freightliner	EAZ	20	2	1	2000-08-12	1	0000-00-00	1
001011190...	2000-10-11	19:05:00	UT6	Patchway	Wales & West	GWZ	90	1	2	0000-00-00	1	0000-00-00	1
001012162...	2000-10-12	16:21:00	NS124	Soho North Jct	EWS	MZ	20	2	1	1998-07-12	1	0000-00-00	2
001012151...	2000-10-12	15:13:00	MP648	Lostock Jct	First North ...	NWZ	20	1	1	0000-00-00	1	0000-00-00	2
001013124...	2000-10-13	12:40:00	BS209	Stansted Mo...	WACN	EAZ	80	1	2	0000-00-00	1	0000-00-00	1
001014152...	2000-10-14	15:29:00	MP182	Manchester ...	Northern Spirit	NWZ	483	6	1	1996-06-02	2	1992-04-15	3
001009001...	2000-10-15	19:21:00	No. 9 Lugton	Lugton	Virgin	ScZ	20	1	1	0000-00-00	1	0000-00-00	2
001015115...	2000-10-15	11:50:00	B12	Alexandra D...	EWS	GWZ	5	1	1	0000-00-00	1	0000-00-00	1
001016140...	2000-10-16	14:04:00	B142	Filton	Wales & West	GWZ	20	7	1	1997-04-28	1	0000-00-00	2

Figure 15.4: Computer-Based Visualisation of SPAD Data

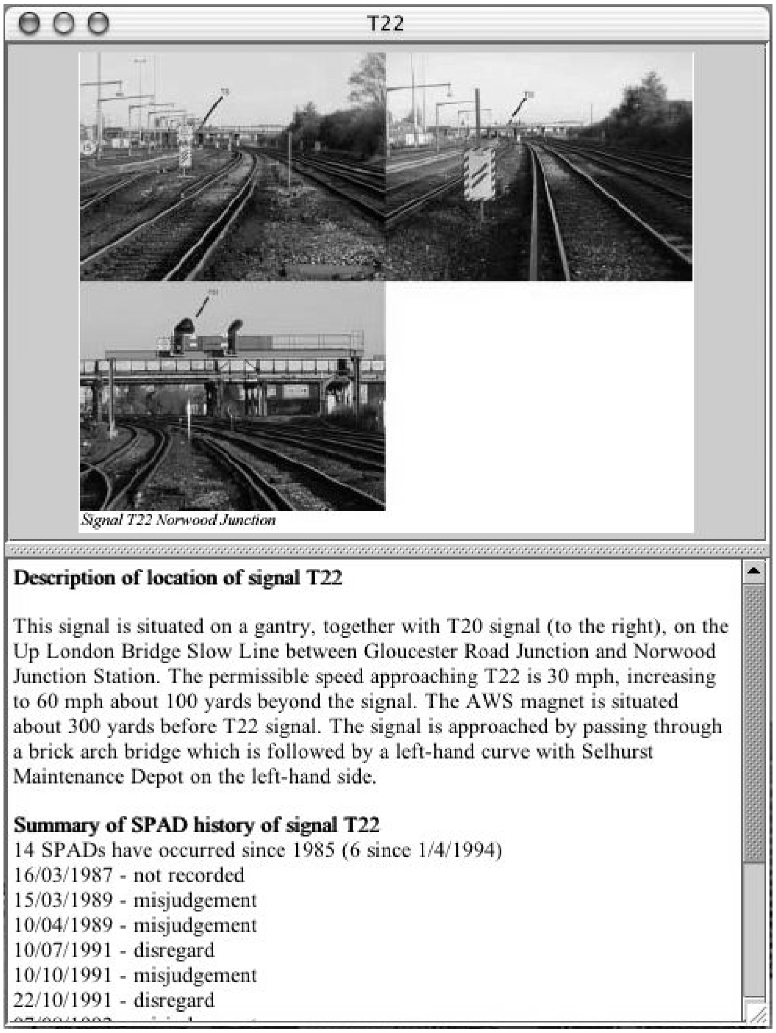

Signal T22 Norwood Junction

Description of location of signal T22

This signal is situated on a gantry, together with T20 signal (to the right), on the Up London Bridge Slow Line between Gloucester Road Junction and Norwood Junction Station. The permissible speed approaching T22 is 30 mph, increasing to 60 mph about 100 yards beyond the signal. The AWS magnet is situated about 300 yards before T22 signal. The signal is approached by passing through a brick arch bridge which is followed by a left-hand curve with Selhurst Maintenance Depot on the left-hand side.

Summary of SPAD history of signal T22
14 SPADs have occurred since 1985 (6 since 1/4/1994)
16/03/1987 - not recorded
15/03/1989 - misjudgement
10/04/1989 - misjudgement
10/07/1991 - disregard
10/10/1991 - misjudgement
22/10/1991 - disregard

Figure 15.5: Signal Detail in SPAD Visualisation

Overshoot slider then they can alter the minimum overshoot distance for any SPAD displayed in the map view on the left. All SPAD incidents with an overshoot that is less than the value displayed on the slider will not be shown as dots on the map. Similarly, if the user selects the right end of the Severity Category slider and moved it from 8 to 3 then only SPAD incidents assessed to be in categories 1, 2 and 3 would be displayed. This form of interaction is known as 'dynamic querying' because it is, typically, built on top of a more conventional database. Each time the user makes a selection and changes the position of the slider, a new query is automatically generated by the visualisation tool and evaluated by the underlying database. It is important to stress, however, that the user is only aware of the interface components shown in Figures 15.4 and 15.6. They need not consider the underlying complexities of relational implementations. The lower portion of Figure 15.6 represents a choice or list widget. The user can filter their query to only display SPADs associated with particular train operating companies. By combining this approach with the slider mechanisms, it is possible to identify the most severe incidents involving particular operators during particular times of the day.

As mentioned, the number of dots shown on the map view in Figure 15.4 is updated in response to each query made by the user. As might be expected, the number of dots shown will be greatly reduced if the sliders are altered so that only the most severe incidents are considered. Conversely,

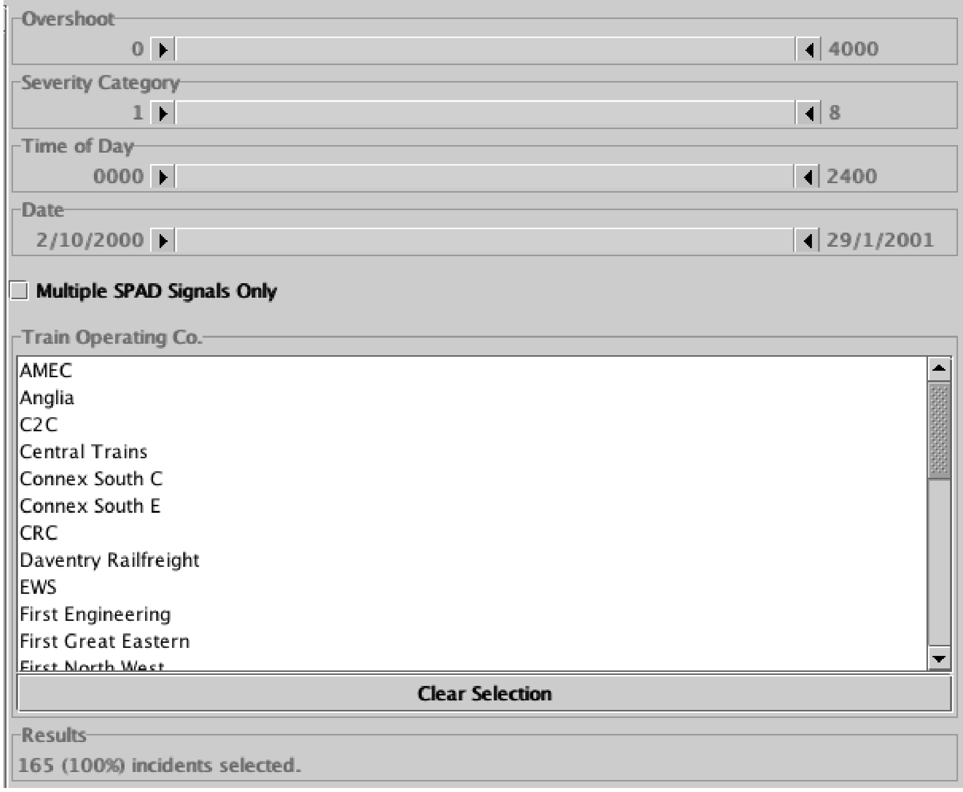

Overshoot
0 4000

Severity Category
1 8

Time of Day
0000 2400

Date
2/10/2000 29/1/2001

□ Multiple SPAD Signals Only

Train Operating Co.

AMEC
Anglia
C2C
Central Trains
Connex South C
Connex South E
CRC
Daventry Railfreight
EWS
First Engineering
First Great Eastern
First North West

Results
165 (100%) incidents selected.

Figure 15.6: Dynamic Queries in the Visualisation of SPAD Incidents

the number of dots will increase greatly if the sliders are then adjusted to consider lower severity incidents. This use of dots on the map view is based on a number of previous visualisations in other domains, including epidemiology, where it is important to monitor the location of particular types of incident [752].

The visualisation, described above, created a number of practical difficulties when applied to incident reporting. For example, many incidents are clustered within a particular geographical location. In epidemiology this can correspond to particular out-breaks of a disease. In our railway case study, these 'hot spots' often corresponded with complex network characteristics and poorly sited signals. This resulted in a large number of dots representing SPADs within a small set of locations on the map. Unfortunately, this representation created problems because users had to 'zoom' in to gain the more detailed map view that was necessary to distinguish between different SPADs on signals that were close together. This was a particular problem given the high density of commuter rail operations in the South East of England and especially around London. Fortunately, other visualisation research can be used to identify potential solutions to this problem. For instance, Figure 15.7 shows how current versions of the prototype exploit a form of Fekete and Plaisant's eccentric labelling technique [246]. As the user moves their mouse over a region in which more than one SPAD has occurred the system will automatically 'pop up' an associated label with the location name of the incident. The user can then select each individual label to gain more detailed information about one of several incidents within the same area.

As mentioned, there are relatively few examples of more advanced visualisation techniques being applied to the monitoring of incident reporting systems. In consequence, the selection of appropriate techniques remains an area of active research. For example, the need to use eccentric labelling emerged as the visualisation prototype was expanded to support larger quantities of SPAD data. This illustrates the way in which the selection of appropriate techniques is driven by the problems that are created by the application domain. Given the lack of research in this area, it is unsurprising

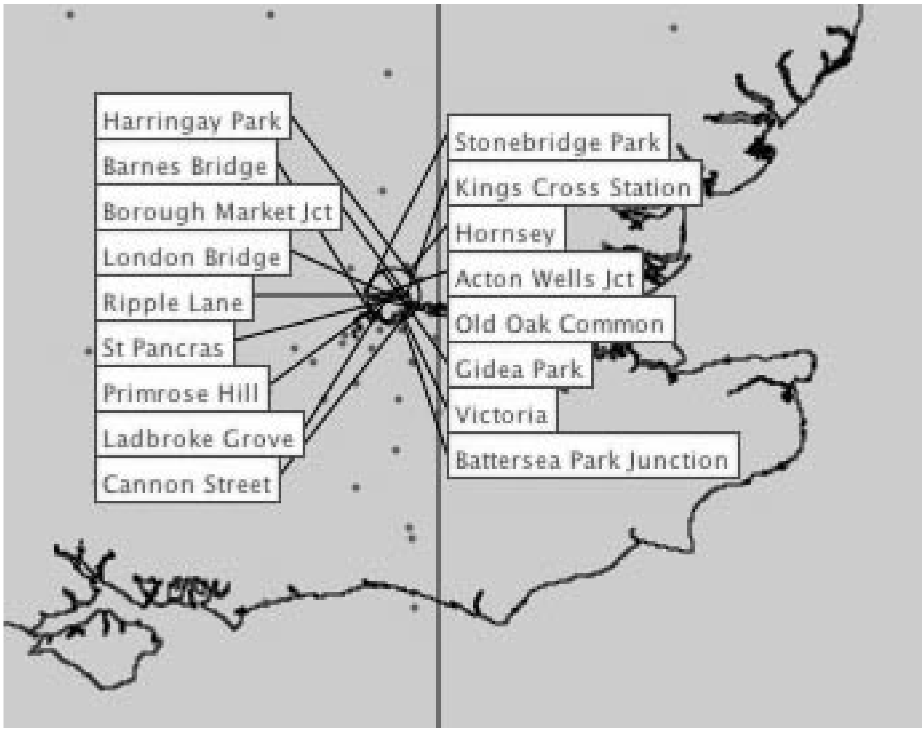

Figure 15.7: Eccentric Labelling in the Visualisation of SPAD Incidents

that the initial attempts described in this chapter should also produce some techniques that failed to support the needs of their intended users. For example, initial discussions with representatives from UK Railway Safety indicated that their safety managers would prefer to integrate the sliders and lists of the dynamic querying technique with more conventional graphs and charts. During 'walk-through' demonstrations, several of their senior managers argued that they were already familiar with the geographical distribution of events and more would be gained by integrating standard forms of statistical representation with the more innovative aspects of dynamic querying. Current research is investigating whether these views are shared by Train Operating Companys. The intention is to extend the initial prototype so that users can alter the image in the top left of the display in Figure 15.4.

Hamming argues that 'the purpose of computing is insight and not numbers' [304]. This section has argued that computer-based visualisations can be used to increase the insights that might otherwise be gained from monitoring statistical summaries of incident reporting metrics. We have not, however, presented any direct evidence to demonstrate the validity of these claims. Previous sections have considered some of the problems that arise from attempts to obtain such evidence. For instance, ethical problems restrict opportunities to introduce new information systems in 'live' reporting schemes if new visualisations can potential hide important information about adverse events. Further problems arise from the novelty of many computer-based visualisations. Users can express strong subjective satisfaction during the initial use of these systems simply because they represent a departure from existing techniques. This initial approval does not, however, imply that they will continue to be satisfied with the system over a longer period of time. The following section, therefore, describes how experimental techniques can be used to gather evidence about the utility of such visualisation tools and about incident reporting schemes in general.

15.4.7 Experimental Studies

The work on SPAD visualisation provides a useful case study of the problems that can arise in assessing or evaluating the meta-level effectiveness of incident monitoring tools. Some of these problems stem from the nature of such information systems. Computer users often find it difficult to express their requirements and needs to a designer. The field of requirements engineering within computer science has developed numerous techniques to address this problem [459]. As has been explained in the previous section, prototype implementations can provide potential users with a better idea of what is possible using existing technology. There is a danger, however, that considerable resources will be invested before safety managers and regulators confirm that the implemented system provides few benefits beyond those of existing systems! Participatory design techniques have been developed to address this problem. End users are represented in design teams and provide detailed guidance on a daily basis. The introduction of this end-user feedback can reduce the likelihood that a final implementation will fail to meet the needs of its potential users. Unfortunately, a number of problems limit the use of such validation techniques for the monitoring systems that support incident reporting schemes. In particular, most systems will only ever be used by a handful of domain experts, safety managers and regulators. These individuals usually play important safety management roles and, hence, their time is both a valuable and scarce resource. In consequence, it is often impossible to secure the level of commitment implied by participatory techniques.

The difficulty of obtaining access to the individuals who are involved in monitoring incident reporting systems creates further problems. As mentioned, key personnel seldom have the time that is necessary to conduct prolonged evaluations of prototype systems. This makes it difficult to perform the longitudinal studies that combat the biases created by the introduction of novel technology and by the Hawthorne effect, mentioned in Chapter 5. Long-term studies also often imply that monitoring systems will run alongside existing applications. This duplication avoids the ethical problems of experimenting with a 'live system'. It also creates additional managerial complexity and considerable expense.

Fortunately, a range of low cost evaluation techniques can be used to address the problems of gaining access to key staff. Some of these methods minimise the direct participation of end-users in the early stages of design. For instance, heuristic evaluations provide designers with rules of thumb that can be used to make inferences about potential usability problems in the final implementation of a computer-based monitoring system. There are both general heuristics [634] and heuristics that support the evaluation of particular systems. For instance, Shneiderman [739] provides a set of criteria that can be used to assess interactive visualisations:

- *overview*: the visualisation must provide the user with a high-level overview of the information that is being presented;

- *zoom*: the visualisation must enable the user to move from the higher-level overview to focus on specific items of interest;

- *filter*: the visualisation must enable the user to filter out related items of information that are not relevant to their current information requirements;

- *details on demand*: it should be possible to select a particular item or group of items and obtain additional information about the selected items;

- *relate*: it should be possible to use the visualisation to view relationships between items of information;

- *history*: it should be possible to undo the effects of a selection or filtering operation, it should also be possible to redo previous operations when necessary.

- *extract*: it should be possible to extract sub-collections from the mass of initial data so that queries can be posed on subsets of the data.

Speirs has used these criteria in the evaluation of the SPAD prototype that was described in the previous section. For example, the dynamic querying facilities provide means of rapidly undoing a filter operation by returning the sliders to their original position. However, early implementation did not enable users to *extract* and save sub-sets of the data for later analysis. It can be argued that this is unnecessary given the ease with which queries can be composed from the simple interface widgets. This argument illustrates both the strengths and weaknesses of heuristic evaluation. These guidelines are a starting point for the evaluation of a potential system. They are also subjective and open to a wide range of interpretations. This makes it likely that different designers may apply the same criteria in a range of different ways. In consequence, there is often a need to perform direct user evaluations to resolve the different claims that can be made about particular heuristics.

Fortunately, a range of low cost techniques can be used to provide user feedback but without the expense associated with longer-term evaluations. For instance, cooperative evaluations and 'think aloud' techniques require that analysts set potential users a series of tasks with a prototype implementation [876]. The focus on accomplishing specific tasks, such as using the system to create a particular statistical summary, avoids the need to ask leading questions, such as 'hwta do you think?'. The responses to such prompts are, typically, impossible to interpret as they can be biased by a range of subjective factors including a concern not to appear ignorant about information technology. By focusing on whether or not the user can perform particular tasks, the intention is to determine whether the system will meet their needs. Subjective satisfaction can be assessed as part of subsequent validation activities. The participants in a cooperative evaluation then attempt to perform the tasks as best they can. If there is any confusion or they do not know what actions to perform then they should express their uncertainty by 'thinking aloud'. The designer can then either provide appropriate feedback or allow the potential user additional time to work on the problem. In either case, there break-downs are noted and become the focus for subsequent re-design. As mentioned, these techniques are relatively low cost because they do not require the prolonged participation of senior staff and domain experts. Feedback is provided in a relatively informal setting and evaluations can be conducted in a relatively short period of time.

Unfortunately, a number of further limitations affect the use of cooperative evaluations to monitor the effectiveness of monitoring tools. In particular, it can be difficult to generalise beyond the results provided for particular users operating a particular version of a prototype implementation. The fact that one safety manager successfully accomplished a task does not imply that others will achieve similar successes. Further problems arise when validation tasks must derive comparative measures for the relative utility and usability of rival designs. It can be difficult to use the introspections derived from 'think alouds' to show that one design is better than another. It is for this reason that several groups have attempted to perform experimental evaluations of monitoring tools [753]. These techniques rely upon established methodologies, often derived from experimental psychology [686]. They rely upon the analysts' ability to distinguish the change, or independent variable, that is linked to a change in the measurement of a dependent variable. For example, an experiment might be conducted to establish the hypothesis that a new monitoring system reduces the time taken to perform a range of key tasks. The dependent variable would be the two versions of the monitoring system. The independent, measured variable would be the timings taken over the range of tasks. The method chosen to conduct such an experiment must be carefully considered to ensure that only the dependent variable is altered between the two conditions. For example, if experienced staff were used with the new system and novice staff with the old then one might expect better performance with the newer system than with the old.

Speirs has used this approach to evaluate the utility of his visualisation tool as a means of monitoring incident reporting data [753]. This study illustrates the complexity of conducting experiments within this area:

1. it is unclear how to develop appropriate tasks for users to perform during the evaluation Traditionally, many validation activities have focussed on well-specified tasks such as finding the answers to particular questions. For instance, a user of the SPAD visualiser might be asked to find out the distance of overshoot associated with an incident in a particular location. This style of experimental evaluation cannot easily be used to support the validation of incident reporting systems. Safety managers are seldom faced with such specific information requests. Greater challenges come from the need to identify patterns and trends within a complex data set. This implies that any evaluation will have to focus on less directed forms of interaction;

2. fatigue can complicate experimental evaluation. The longer that a user interacts with a tool then the more tired they can become, especially if they are being asked to use new and unfamiliar systems. One consequence of this is that if an evaluation compares two systems then many users' will perform less well with the second system that they meet. Fatigue becomes more of an influence than any potential design improvements. Counter-balancing can be used to address this problem. This implies that half of the user group will meet the old system first and the other half will meet the new version first;

3. learning effects with new systems complicate experimental evaluation. Counter-balancing cannot reduce the problems created by learning effects. For instance, it may take some time before novice users of a new system can build up the same level of expertise that they have achieved with the existing implementation. This effect can often be observed when users' performance with a new system slowly improves as they attempt successive tasks. This problem can be addressed by designing a series of training tasks to ensure that users are happy with both of the systems that are being compared. The users' performance with these tasks is not measured as part of the experimental evaluation and users are only encouraged to progress once they feel happy that they are able to perform them unassisted;

4. learning effects with particular tasks complicate experimental evaluation. Learning effects not only complicate the comparison of alternative monitoring systems. They also effect the tasks and the data sets that are presented to the user during the evaluation. For example, if users were asked to perform the same tasks with two different systems then one might anticipate that a knowledge of their previous answers might influence subsequent responses. It is for this reason that many experimental evaluations provide different tasks for each system. This creates further problems because some questions might be 'easier' than others. Hence, it becomes necessary to ensure that counter-balancing also considers the questions that are associated with each system. Table 15.16 provides an example of the complexity that this can create. It also emphasises the point that was made earlier, experimental evaluations can require access to relatively large numbers of users in order to exploit such techniques;

Group 1	Old System	Questions A	Questions B	New system	Questions A	Questions B
Group 2	New System	Questions A	Questions B	Old system	Questions A	Questions B
Group 3	Old System	Questions B	Questions A	New system	Questions B	Questions A
Group 4	New System	Questions B	Questions A	Old system	Questions B	Questions A

Table 15.16: Counter-balancing Systems and Tasks

5. learning effects with particular data sets complicate experimental evaluation. It can be difficult to use counter-balancing as a means of reducing learning effects associated with particular data sets. For example, users may get a better 'feel' for the information that they are being asked to monitor as they interact with it over time. This effect could be addressed by partitioning the incident dataset into two or more sections and then introducing additional user groups to experience each data set with one of the experimental conditions, shown in Table 15.16. Such techniques undermine the *validity* of the evaluation. It is likely that any final implementation would have to support the entire available data set. Frequently the need to control experimental conditions can lead evaluators to impose unrealistic constraints on the use of a system. This creates problems because the results of any study are, therefore, indicative of the system that was evaluated and not of the system as it might operate in an eventual working environment;

6. what do we measure? It can be difficult to identify the measures that might be used to assess the effectiveness of a monitoring system. It is seldom the case that fine grained timings would have a significant impact upon many safety managers. It may make little different whether it takes 6 or 8 minutes on average to identify a particular trend so long as that trend can be identified. Further problems arise because measurement criteria are, typically,

multi-dimensional. In consequence, it can make little sense to compare very different systems using the same criteria. For instance, relational databases can provide relatively fast access to information in response to specific queries. In contrast, the SPAD visualisation tool supports less directed forms of search. It would be unsurprising to find that each tool performed less well when assessed against criteria that were not used to guide their initial development.

This is a partial list, many further problems complicate the design of experimental evaluations. The interested reader is directed to the summary in [127, 303]. In contrast this paragraph provides an example of the experimental techniques that were used in the validation of the initial SPAD prototype.

It was initially only possible to recruit seven subjects. This illustrates the way in which access constraints can frustrate attempts to employ the counter-balancing mentioned above. As this was a preliminary evaluation, Spiers compared the performance of the SPAD visualisation tool with the statistical presentation of SPAD data that is presently hosted on the HMRI's web cite [354]. Further studies are currently comparing the visualisation tool more directly with the existing presentation of data from Railway Safety's extended Safety Management Information System (SMIS), mentioned in previous sections. The initial comparison focussed on a number of relatively open ended tasks. Users were asked to rate their agreement with a number of statements on a scale from one to seven [753]. As we shall see, this complicated the analysis of the results from the study. It was, however, intended to provide a measure of the certainty that participants felt in the conclusions that they reached using a particular visualisation. For the purposes of counterbalancing, the questions were assigned to one of two groups and the position of questions within each group was varied. From this it follows that each participant answered all of the questions, however, the order that they answered them was varied as was the system used to generate their answer:

- Set A:
  1. Events at multiple SPAD signals constitute around 40% of all SPADs.
  2. Most SPADs have a severity category of 3 or over.
  3. Most SPADs involve an overshoot of less than 200 yards.
  4. The number of incidents at multiple SPAD signals is increasing from month to month.
  5. Incidents involving a signal that had previously been passed at danger usually also involve a driver that has been involved in a previous SPAD.
- Set B:
  1. Railtrack Midland Zone (MZ) is the zone with the lowest number of SPADs.
  2. The number of incidents is relatively stable from month to month.
  3. SPADs are more common in the morning (24.00-12.00) than in the evening (12.00-24.00).
  4. The incident with the longest overshoot distance occurred in Manchester.
  5. No SPAD has occurred north of the incident at Perth Yard (signal P197).

Some questions can be answered directly from the information provided by a particular system. For instance, the map view of the visualisation can be used to directly identify whether or not a SPAD occurred beyond signal P197. Similarly, the spreadsheet view provided by the HMRI, illustrated by Table 15.15, can be directly used to identify whether or not the longest overshoot occurred near Manchester. Other questions involved a greater degree of analysis and interpretation. For example, there is no direct means of determining whether multiple SPADs formed a particular percentage over overall incidents from the data that was presented to each user in either system.

The degree of interpretation and analysis involved in some questions, together with the seven point scale, created problems in analysing the results of the evaluation. In order to do this, ideal responses were identified for each question. These 'solutions' were based on the HMRI's interpretation of the SPAD data. The users' performance with each system was measured in terms of the absolute (ie., non-negative) divergence of their score from this ideal value. For example, a user might assign the value 1 to show that they disagree with the statement ‘most SPADs involve an overshoot of less than 200 yards’. If the HMRI indicated that most SPADs did involve overshoots of less than 200 yards then the ideal score would have been 7. The user would then be assigned the value 6 for their performance on this question to indicate variance from the ideal answer. The results from this evaluation are presented in Table 15.17. As can be seen, timings were not taken during the study and users were encouraged to take as much time as they liked.

Subject	SPAD Visualisation	HMRI Site
A	9	9
B	6	9
C	6	4
D	4	10
E	13	1
F	2	5
G	5	10
Total variance from ‘ideal’	45	48

Table 15.17: Initial Results from Experimental Evaluation of the SPAD Browser

As can be seen, the results showed very little difference between the performance of the users with either visualisation. Perhaps more remarkable is the variation in individual performance. For example, subjects A and B did equally well with either monitoring system. In contrast, subjects D did much better with the SPAD visualisation while E showed less variance from the ‘ideal’ responses when using the more conventional spreadsheet format used by the HMRI. These results led Speirs to realise that users were exploiting the graphical visualisations and the tabular format for different purposes. The sliders of dynamic querying techniques were used to filter the data set while the tabular or spreadsheet view was used to rapidly scan for particular numeric values. This justified the introduction of tabular data into the bottom half of the display illustrated in Figure 15.4.

Subsequent work identified a number of limitations with this experiment. These limitations illustrate further problems with the experimental method as a means of validating incident monitoring tools. Cooperative evaluations were conducted with more senior staff from Railway Safety. They argued that the tasks were not particularly significant for the end-users of SPAD incident data. The location information provided by the map view simply helped to reinforce the correlation between SPADs and the density of railway operations within particular zones. In contrast, they advocated the geographical presentation of information about suicides as well as trespass and vandalism incidents. They argued that these incidents did not relate so directly to the flow of traffic but did possibly cluster around particular geographical regions, for instance with particular problems of social deprivation and unemployment [753].

The previous paragraphs have explained how ethical issues can prevent investigators from evaluating new reporting techniques on ‘live’ systems. Initial ‘teething’ troubles can lose safety-related information. Dissatisfaction with prototype tools and techniques might jeopardise the future success of an existing scheme. Laboratory techniques, typically, avoid these problems by examining the operation of a new system in under carefully controlled conditions with simulated tasks. Some attempts have, however, been made to conduct experimental comparisons with ‘live’ systems. These evaluations have abandoned some of the controls that are normally associated with laboratory assessments in order to increase the experimental validity of their results. Novel techniques are compared to existing approaches using real operators reporting actual incidents. This approach was used to evaluate initial versions of the ATSB’s INDICATE program [46]. INDICATE provides company’s with a framework for eliciting, documenting and monitoring safety-related incidents. It is also supported by a range of software tools that can be tailored to the individual requirements of participating organisations. The INDICATE program has also been extended to support organisations in the aviation, road, rail and maritime industries. Initially, however, the evaluation focussed on the aviation domain. An Australian regional airline agreed to use INDICATE in one of its operational bases while another section of the same organisation was used as a control group over an eight month trial period The length of the evaluation reflects the intention to recreate 'realistic' reporting behaviours during the study. Experiments that run over a couple of hours can often suffer from biases that create 'atypical' reporting behaviour. This style of evaluation has much in common with the Sentinel studies, described in previous paragraphs. Resources are focussed on an initial trial of a new reporting technique. There are some important differences. In particular, the use of a control group is directly taken from experimental evaluation techniques and helps to provide some basis for comparing the results obtained by the introduction of a new scheme. As we shall see, however, it can be difficult to make accurate comparisons between these two different groups.

Five evaluation criteria were used to determine whether INDICATE had a positive effect: airline safety culture; staff risk perception of safety hazards; willingness of staff to report safety hazards; action taken on identified safety hazards and staff comments about safety management within the airline. 48 safety reporting forms were submitted by the INDICATE group, 9 were submitted by the Non-INDICATE group. Analysts argued that this difference 'may be a direct result of an attitude change within the INDICATE base as a result of the safety program (e.g. a more positive attitude to reporting safety issues, increased staff confidence that safety problems would be addressed, more awareness amongst staff of operational hazards, and improved staff commitment to improving company safety)' [46]. A questionnaire was used to provide feedback about the other evaluation criteria. A reliability analysis was conducted to establish that the questions elicited consistent responses from individual participants. Such studies are importance because doubts can arise if the same person provides radically different answers to the same questions within a short period of time or if individuals in the same organisation express radically different views about the same topic. Under such circumstances, it may be more important to understand the differences within a group than between two different groups. After the trial, the INDICATE group showed 'a significant improvement in their mean safety culture score, while the Non-INDICATE base results showed a poorer safety culture'. Various statistical measures, including T-tests and ANOVAs, were used to support the finding. It was argued that there was a '99.9% certainty that the safety culture improvement, demonstrated in the INDICATE base, was due to the implementation of the safety program and not some chance factor' [46]. Such arguments are, however, complicated by the problems of conducting 'experimental' evaluations on reporting systems. At the start of the study, staff in both centres indicated that safety was managed 'in a positive manner' [46]. However, the INDICATE site achieved a slightly better initial score than the control group. This difference makes it difficult to interpret the results, cited above. Any subsequent change in attitudes or reporting behaviour can be explained in terms of the initial differences between the two groups. The initial score of the INDICATE group may reflect a pre-existing increase in awareness about safety issues. The subsequent improvement might, therefore, be part of this previous trend rather than a 'direct result' from introducing the new reporting programme.

There was a smaller difference between the two groups in the risk perception questions from the initial questionnaire. At the end of the evaluation period, however, the INDICATE group showed a significant decrease in the risks that it associated with particular hazards. The non-INDICATE group exhibited a much smaller reduction. These results were not expected. The analysts argued that they might have been due to chance factors that did not affect the other metrics. Alternatively, the reduction in the level of risk perceived by the INDICATE group might show that the program had provided staff with a clearer idea of the hazards facing their industry. The reduction in risk perceived by the non-INDICATE group was argued not to be statistically significant. Again, however, these arguments were complicated by the difficulty of identifying the direct influence of the INDICATE program.

It was argued that comments from staff in the INDICATE group revealed that there was: 'better provision of safety training to new staff; more management praise for safe working; better company feedback regarding safety performance; and an increased frequency of safety audits' [46]. These comments were also interpreted to show that staff in the INDICATE group were more confident in their safety management systems. Staff in the non-INDICATE group were 'generally more negative' about communication from management and the reporting of safety incidents. The qualitative assessment criteria were not, however, explicitly documented. This is important because it can be difficult to agree upon the best means of extracting such conclusions from the informal comments of individual workers. For instance, respondent validation can be used to ensure that staff within the INDICATE group agreed with the summary of their comments. Similarly, staff within the non-INDICATE group might have been asked to check the interpretation of their comments. This method can be difficult to apply when comparisons are made between two groups. Unless respondents have access to the comments of their colleagues, they cannot judge whether their responses were 'more negative' than those of another group. Alternatively, independent assessors can be used to summarise respondent comments without knowing the experimental context in which they were obtained. The form of 'blind' reviewing can, however, be difficult if analysts lack the information that is necessary to interpret particular comments. In either case, the key point is that some form of additional validation is often required to support the interpretation of qualitative responses.

A slight variation on the use of experimental techniques in 'real world' settings is to conduct limited studies to support the gradual introduction of a regional or national system. For example, in October 1998 the FRA awarded a 3-year contract to design, develop and test a Toll-Free Emergency Notification System (ENS). This was intended to centralise the reporting of problems at highway-rail intersections [240]. The ENS System was initially installed along limited areas of track within the State of Texas. Early in the project, a number of liability issues were identified and special legislation was required to authorise particular organisations to manage such a facility. The initial study was extended from Texas into Connecticut and then to areas of Pennsylvania. The Pennsylvania program modified the Texas system so that it could be operated in an existing 911 emergency centre. It began by supporting eight selected railroads but the longer-term objective was to 'continue refinement, based on operating experience with the demonstration system, so that a system suitable for statewide usage by short-line railroads is realised' [240].

The ENS project illustrates the way in which a Sentinel-style approach can be integrated into a form of iterative development. The ENS project also considered a number of issues that complicate the use of experimental techniques with 'live' reporting systems. For example, the sample had to broad enough to make it likely that incidents would be reported. It was equally important, however, that the size of the study was not so large that it overwhelmed the available resources. Signs had to be deployed at all public railway crossings along the eight chosen railroads. They also had to be deployed a private crossings that were considered active enough to create a significant risk of a potential incident. Farm field crossings that were used two or three times per year did not warrant a sign. This may have prevented the reporting of some incidents but also helped to focus the allocation of finite resources. The ENS evaluation illustrates further problems that complicate the use of experimental methods on 'live' systems. In particular, there was a concern that sufficient data should be available about the safety record of the existing system. Without this data, accurate comparisons could not have been made following the introduction of the new ENS application. Unfortunately, the under-reporting of adverse events makes it very difficult to obtain accurate data. There can also be a range of more prosaic problems. The ENS analysts also had to ensure that they obtained accurate information about the location of the existing crossings.

A similar approach to that adopted in the ENS study was also used to examine a range of techniques that were intended to reduce the use of train horns at railway crossings in the United States [235]. This, in turn, was intended to reduce environmental problems associated with the use of train horns to warn drivers and pedestrians of an approaching train. This study is interesting because it mixes elements of several different monitoring techniques. An experimental method is used in that the study attempted to control conditions around a number of road-rail crossings so that comparisons could be made between driver behaviour with different protection mechanisms. The techniques also borrowed from the observational approach mentioned in previous sections because video taping was used to record the behaviour of 'real' drivers as they approached the crossing. The results of this analysis not only provided insights into the effectiveness of safety measures that were intended to address previous incidents. The video analysis also provided information about the reliability of reporting systems because it helped to identify a range of 'near miss' incidents in which accidents were narrowly avoided but which would not otherwise have been notified to regulatory organisations or the operating companies.

Previous paragraphs have looked at the application of experimental techniques to support the meta-level monitoring of incident reporting systems. In other words, we have concentrated on the evaluation of innovative reporting techniques rather that on the performance of an individual reporting system. The same evaluation methods can also be applied more directly to anticipate whether the nature of incidents will change as a result of revised operating procedures. The study involved an agreement between Spokane County, the Washington State Utilities and Transportation Commission and the Burlington Northern Sante Fe Railroad. It was based around four phases. This helped to ensure that the behaviours, which were witnessed during the study period, provided accurate insights into longer-term driver performance. The first phase provided a 'control' or 'baseline period'. During this time there was neither a median barrier nor a whistle ban. The second phase of the experiment introduced a barrier but did not enforce a whistle ban. The third phase involved the introduction of median barriers and a whistle ban. Each of the first three phases lasted 115 days each. The final phase monitored driver behaviour for one week in each of the following three months and one week each quarter for the year after the original study. During this time, the median barriers remained in place together with the whistle ban.

The analysts focussed their attention on incidents in which vehicles and pedestrians continued to cross even though the crossing had been activated. An incident was also defined to have occurred if a pedestrian or vehicle collided with the gate or if they went around a gate after it had been activated. They also argued that most attention should be devoted to those incidents where there was a train present. Their study identified numerous cases in which the gate activated without a train being present, for instance through gate malfunction. Table 15.18 summarises incident frequencies for each of the four phases in the study. There was a sharp decline in the incident rate between phase 1 and 2. The percentage of gate activations in which there was an incident fell from 34% to 1.2% after the introduction of the barriers. There was little change after the introduction of the horn ban between phases 2 and 3. There was a relatively small increase in incidents during the final phase.

	Phase 1	Phase 2	Phase 3	Phase 4
No. of Gate Activations				
Train present	4,556	4,924	5,003	680
No train present	31	155	117	18
No. of Incidents				
Train present	1,565	61	66	14
No train present	419	5	9	7

Table 15.18: Number of Gate Activations and Incidents, With and Without a Train Present

This mix of experimental techniques in 'real' contexts with more observational techniques illustrates a range of further practical problems. Residents reported frequent soundings of the train horn even during phases three and four when the bans was in place. A dosimeter placed at this location expressly to monitor possible whistle soundings revealed 231 noise spikes in Phase 3 alone. As in previous studies, it was also important to ensure that accurate comparisons could be made between the results for each of the different conditions. In this case, it was necessary to ensure that there were no differences in the number of times that the crossing gates were activated between the different phases of the project. In Phase 1, there were 4,587 activations in all at an average of 39.9 per day. In Phases 2 and 3, the averages were slightly higher 44.4 and 44.8 respectively. The shorter periods of Phase 4, yielded an average of 38.8. Unfortunately, the relative similarity in the number of gate activations did not characterise the statistics for car/automotive traffic. The average annual daily traffic for the first phase was estimated at 3,831 cars per day. This was significantly higher than the 1,918 cars in phase 2 and the 1,991 in phase 3. This statistic was not calculated for the more limited observations in period 4.

These problems illustrate the difficulty of accounting for the many different variables that might influence the incident rate between a number of 'experimental conditions'. It is difficult to know how to interpret the results from such a study. It might be argued that the relatively high number of automobile journeys observed in the first period invalidates the use of the incident data for the control of baseline phase. It is important to emphasise that those involved in conducting the experiment would almost certainly have been unaware of the potential imbalance as they conducted the study. It would only have become apparent during the detailed analysis of the video data. Given such objections, analysis would have been forced to remove the barriers and the whistle ban to recreate the conditions under which the first phase was conducted. However, there would then be the problem of ensuring compliance with staff who had already become familiar with the ban. The ethical implications of removing barriers must also be considered. The FDA analysts, therefore, simply present the data and identify the potential flaw in the evaluation. This pragmatic stance enables individuals to form their own judgements about the validity of the trial.

In passing, it is important to note that this study identified a number of important insights about driver behaviour at railway crossings. Many of these insights could not be obtained through more conventional reporting systems, For instance, many cars and pedestrians crossed after the gates had been activated but in situations where it was clear that no train was present. For instance, the seven incidents in phase four of table 15.18 represent seven cars crossing one after the other when the gates had failed. The ways in which such events might have led to more serious violations were not considered as part of the study. However, such observations illustrate the way in which monitoring techniques can produce a broad range of insights into the limitations of reporting systems. Until the study took place, nobody reported these ‘successful violations’ in which users were forced to cross an activated crossing in order to cope with a gate failure.

There are further examples of experimental techniques being used to support the monitoring of incident reporting systems. For instance, the FRA’s Volpe Center often performs empirical studies to validate the recommendations that are made in the aftermath of adverse events. Their Transportation Technology Center (TTC) was specifically designed to test all types of rail equipment and vehicles in a variety of weather and terrain conditions. Their facilities allow a limited form of replication in which collisions can be recreated. During these tests certain factors are kept constant while others are systematically varied to support particular hypotheses, in the manner described in previous sections. For instance, identical rolling stock can be operated at different speeds along the same track. These full-scale crash simulation are increasingly used to validate computer models, such as the finite element simulations that are needed to describe the nonlinear properties of material behaviour in rail collisions. These models take the place of destructive testing [851]. There is a particular sense in which these activities help to validate the products of the FRA’s reporting systems. The Volpe center conducts crash testing and similar experimental studies to help determine which adverse events pose the greatest concern for the future safety of the railways. The results of their work should identify those conditions that are known to pose the greatest threat from accident reports. They may also validate the potential for future accidents by identifying problems that have not yet resulted in injuries or fatalities. The results of these experiments help to determine the potential consequences of events described in ‘near miss’ incident reports. This integration of experimental testing and the reporting of adverse events illustrates the way in which many safety-critical organisations are addressing the problems created by causal asymmetries. In the aftermath of an adverse event, we often cannot be sure which of several possible causal ‘paths’ actually led to an observed outcome. Full-scale simulations and computer models can be used to recreate the circumstances leading to an incident or accident. This increases confidence that a causal explanation can account for the observed event .

15.5 Summary

This chapter has identified several different forms of monitoring that can be used to assess the utility of an incident reporting system. For example, accident rates can be used to assess the impact that the recommendations from these schemes have upon the safety of application processes. In particular, analysts can look for evidence that training and operational practices have been directly improved by the insights from reporting systems. They can also demonstrate the effectiveness of such schemes by assessing the contribution that they make to the calculation of future failure rates and consequence assessments during subsequent risk assessments. These validation activities focus directly on the impact that a reporting system has upon the safety of underlying applications.

In contrast to these outcome measures, other forms of assessment focus on the processes that are used to derive recommendations from incident reports. The success of a reporting system can be assessed in terms of the number of reports that are elicited from staff and management. Alternatively, analysts might focus on the efficiency of a scheme by monitoring the costs associated with analysing each report. In particular, process metrics can be devised to calculate the percentage of contributions that result in particular safety measures being introduced into the 'target' organisation.

Incident reporting systems offer a number of additional benefits beyond the specific recommendations that are made in the aftermath of an adverse event. In particular, participation in these schemes can also have a more general effect in raising awareness about safety-related issues. A third set of monitoring techniques, therefore, focus less on outcomes or on process metrics and, instead, focus on acceptance measures. Analysts can assess the effectiveness of a reporting system in terms of the contribution that they make to a wider 'safety culture'. They can also determine whether potential contributors are satisfied that a system is both confidential and unbiased.

Numerous problems complicate the monitoring of incident reporting systems using these three different approaches. Outcome measures are difficult to gather and hard to interpret. In many industries, there are significant concerns about the accuracy of accident statistics. Incident reporting systems also, often, form part of a wider range of measures that are intended to improve the safety of application processes. This makes it difficult to demonstrate that any changes in accident rates are directly due to the introduction of a reporting scheme. At a lower level, recommendations to change employee training and operational practices in the aftermath of an adverse event can be identified as specific benefits from a reporting system. Often, however, these changes can introduce new failure modes. The impact of such changes can, therefore, only be assessed over a relatively prolonged timescale.

Similar problems frustrate the use of process metrics to support the monitoring of incident reporting systems. For example, a paradox of many reporting systems is that the introduction of such schemes will typically increase rather than reduce the number of reported failures. This has led some safety managers to focus on the criticality of reported incidents rather than on submission rates as a metric to measure the impact of such schemes. If a reporting system is having a positive effect then the number of submissions should remain high but the potential consequences of any adverse events should decline as necessary interventions are made in response to previous reports. This approach is difficult to apply effectively because it requires a relatively sophisticated means of measuring the potential consequences of an adverse event. Investigators frequently show considerable disagreement over the potential outcome from the same adverse event. Other process measures suffer from similar problems of subjective interpretation. Attempts to monitor the performance of individual investigators are complicated by the need to determine a 'gold standard' for causal analysis and the generation of recommendations. Attempts to measure the proportion of submissions that lead to safety interventions can be rendered ineffective when a high number of relatively unimportant recommendations might be considered to have a greater impact than a smaller number of far-reaching innovations.

It can also be difficult to use acceptance measures to support the monitoring of reporting systems. It is far from easy to agree upon a set of metrics that can provide adequate feedback about the impact of such schemes on the safety culture or climate within heterogenous organisations. Similarly, different individuals within the same working group can have radically different opinions about the probity or 'trustworthiness' of a reporting system. In such circumstances, monitoring systems can help to identify the diversity of views but provide little help in encouraging greater confidence.

The second half of this chapter has reviewed a range of monitoring techniques that are intended to address some of the problems that complicate the use of outcome, process and acceptance metrics. For example, public hearings, focus groups, working parties and standing committees all provide means of monitoring incident reporting systems. Focus groups can help regulators and safety managers to assess attitudes towards these schemes within sections of the workforce. In contrast, standing committees provide a more sustained framework that can be used to coordinate a range of monitoring activities over longer periods of time.

Subsequent sections considered ways in which incident sampling can be used to focus monitoring activities. One approach is to select a random sample of reports so that they can be followed as they are analysed and recommendations are implemented. Alternatively, monitoring activities can be focussed on the response to particular types of report. Resources might be allocated to see if there are any problems in the way in which reports from particular user groups are handled. For instance, the Ladbroke Grove enquiry focussed attention on the way in which SPAD reports were handled within the UK rail industry. Such sampling techniques create considerable methodological problems. If employees become aware that attention is being paid to particular incidents then it can become more likely that these adverse events will be reported and analysed in greater detail than might otherwise be the case. This illustrates another example of the Hawthorne effect, introduced in Chapter 5. Individuals will alter their patterns of behaviour if they know that their actions are being observed.

The Hawthorne effect that complicates the interpretation of insights gained from incident samples is actively exploited in Sentinel monitoring systems. This approach deliberately sensitises particular groups or organisations so that they are more likely to report adverse events. These groups are given additional training and resources that could not be provided throughout a mass reporting system. The incidents that are reported by the participants in a Sentinel system can then be compared to reporting patterns throughout an industry. This approach can provide insights into the under-reporting problems that affect many schemes. It can also be used to conduct limited evaluations of additional training materials that might eventually be distributed more widely throughout a reporting scheme. As mentioned, Sentinel systems do not overcome the problems of the Hawthorne effect. Participants are, typically, aware that their reporting behaviour is being monitored. Sentinel systems also rely upon sensitising employees so that they are more aware of the adverse events that should be reported. This can have the paradoxical effect of reducing the likelihood of these events. Individuals are less likely to be involved in particular incidents if they have already been warned about them in their Sentinel training.

Observational techniques avoid many of the biases that can be introduced through incident sampling and Sentinel schemes. These approaches rely upon investigators conducting detailed studies of the reporting behaviours in particular working groups. Such techniques often reject the suppositions of theoretical frameworks that can bias the subsequent interpretation of any observations. The intention is to let particular theories about reporting behaviour develop in a bottom-up way. Theories are grounded in the observations rather than confirmed by them. Unfortunately, a number of methodological problems complicate the application of these approaches. In particular, the relatively low frequence of adverse events implies that observers may have to spend many months studying a particular group of employees before they witness any 'reporting behaviour'. Observational techniques are resource intensive. It can take several hours to analyse a single hour of video tape . There are also problems with generalising from the insights gained by monitoring a particular work group. One team's reporting behaviour can provide relatively few insights into the reporting behaviour of their colleagues in different companys, different geographical regions etc [303]. The rich qualitative data that is derived from these studies cannot easily support the statistical analyses that are often required by governmental and regulatory organisations. For instance, many industries are required to monitor their reporting systems as part of a wider assessment of their safety management schemes. In such cases, normalised accident statistics are used to provide direct outcome metrics. Unfortunately, it can be difficult to identify appropriate normalising factors that reflect the diversity of many industries. For instance, the safety of a rail operator might be assessed in terms of the number of fatalities per passenger mile. Such a normalisation could not easily be applied to freight operations. It might also provide few insights for networks in which marshalling operations accounted for the majority of fatalities.

Previous sections have reviewed a range of additional factors that complicate the use of statistical techniques to support the monitoring of incident reporting systems. In particular, it can be difficult for managers, regulators and the general public to correctly interpret the values that are derived from more sophisticated forms of statistical analysis. It is for this reason that many incident reporting systems are being supported by computer-based monitoring systems. Visualisation tools enable managers to explore and exploit a range of statistical information about their schemes. We have illustrated this argument by describing a SPAD tool that provides railway regulators and operating companies with ‘dynamic querying’ techniques.

The chapter concluded by arguing that a form of meta-level monitoring must be conducted to ensure that visualisation systems and similar monitoring tools actually support their intended users. In the case of the SPAD tool, mentioned above, the use of a geographical information system to provide feedback on the location of incidents was perceived to be less important that the provision of statistical information in the form of more conventional charts and graphs. Laboratory-based evaluations provide, arguably, the best developed set of methods for validating such meta-level monitoring tools. These techniques rest upon experimental situations in which it is possible to distinguish the change, or independent variable, that is linked to a measure of the dependent variable. Unfortunately, the ecological validity of laboratory-based experiments is often questioned. In other words, the controls that are necessary to isolate the dependent variable often creates situations that are a long way from those that characterise the working lives of incident investigators. For instance, they may be required to use monitoring tools in purpose-built evaluation labs. These, typically, exclude sources of distraction including their colleagues, telephones broken printers etc. Fortunately many of these problems are being addressed by hybrid techniques that combine elements of observational and experimental validation. For instance, the FRA have used video cameras to record driver behaviour at rail-road intersections. These observational techniques have been used to monitor the safety improvements provided by a range of barriers and warnings. These different measures are varied in a systematic way that borrows from the use of control groups and different experimental conditions in the laboratory studies mentioned above. Although such techniques overcome some of the objections that have been made towards both laboratory-based studies and observational techniques, they also introduce a range of ethical concerns. There is a danger that lives will be lost in control groups or experimental conditions that are deliberately deprived of certain safety features.

Chapter 16

Conclusions

This book provides a broad survey of incident reporting techniques. This focus is justified because national and international reporting schemes have recently been established to support the aviation industry [308], chemical production [162], marine transportation [387], military acquisition [287] and operations [806], nuclear power production [382], the rail industries [664], healthcare [105]. Chapter 2 has presented a wide range of reasons why incident reporting systems are being used to support the operation of safety-critical applications. For example, incident reports help to find out why accidents do not occur. Many incident reporting forms identify the barriers that prevent adverse situations from developing into a major accident. These insights are very important. They help analysts to identify where additional support is required in order to guarantee the future benefits of those safeguards. The higher frequency of incidents permits quantitative analysis. It can be argued that many accidents stem from atypical situations. They, therefore, provide relatively little information about the nature of future failures. In contrast, the higher frequency of incidents provides greater insights into the relative proportions of particular classes of human 'error', systems failure, regulatory weakness etc. Incident reporting systems also provide a reminder of hazards. Incident reports provide a means of monitoring potential problems as they recur during the lifetime of an application. The documentation of these problems increases the likelihood that recurrent failures will be noticed and acted upon. Reporting systems can also provide feedback that keeps staff 'in the loop'. Incident reporting schemes provide a means of encouraging staff participation in safety improvement. In a well-run system, they can see that their concerns are treated seriously and are acted upon by the organisation. Greater insight into national and global safety issues can be gained.

A further argument in favour of incident reporting schemes is that they can encourage the sharing of safety-related lessons and data. Incident reporting systems provide the raw data for comparisons both within and between industries. If common causes of incidents can be observed then, it is argued, common solutions can be found. However, in practice, the lack of national and international standards for incident reporting prevents designers and managers from gaining a clear view of the relative priorities of such safety improvements. Incident reporting schemes are also cheaper than the costs of an accident. The relatively low costs of managing an incident reporting scheme should be offset against the costs of failing to prevent an accident . This is a persuasive argument. However, there is also a concern that punitive damages may be levied if an organisation fails to act upon the causes of an incident that subsequently contribute towards an accident.

Many of these potential benefits have been cited by the politicians and civil servants who are responsible for setting up new schemes. On the 13th June 2000, the UK Health Secretary, Alan Milburn, and the Chief Medical Officer for England, Liam Donaldson, announced the establishment of a centralised reporting facility for adverse incidents across the UK National Health Service (NHS) [105]. The Chief Medical Officer said; "At the moment there is no way of knowing whether the lessons learned from an incident in one part of the NHS are properly shared with the whole health service". The Health Secretary said; "Patients, staff and the public have the right to expect the NHS to learn from its mistakes so we can ensure the alarm bells ring when there are genuine concerns so they can be nipped in the bud". On April 2nd, 2001 the U.S. Department of Transportation announced a reduction in the threshold for reporting hazardous liquid pipeline incidents from 2,100 gallons, or 50 barrels, to just 5 gallons. This extension of the reporting system was intended to 'heighten the quantity, quality, and usefulness of reported accident information' [704]. In 2002, the FAA Administrator Jane Garvey stated that the GAIN international incident reporting initiative "is one of our best hopes for enhancing aviation safety in the next century". David Hinson, the Former FAA Administrator, observed that "the inception of the GAIN concept" was the most "significant accomplishment of my tenure" [656].

Numerous studies have described tools and techniques that are intended to help realise the potential benefits of incident reporting systems [444, 843]. Chapter 15 has, however, argued that very few papers analyse the reasons why previous initiatives have failed to yield these safety improvements [409]. This is a significant omission. Some reporting systems only elicit a very small number of contributions. The Royal College of Anaesthetist's recently concluded that self-reporting only retrieves approximately 30% of the incidents that are detected by independent audit [715]. Those submissions that are obtained may only come from particular sections of the workforce [119]. Under-reporting is, however, only one of the problems that complicate the introduction and maintenance of incident reporting systems. For example, finite resources can also be exhausted if schemes elicit too many contributions. The difficult balancing act between encouraging participation and discouraging excess contributions can be illustrated by the FDA's MedWatch program. Clinicians do not have to be convinced that a device has actually caused an incident in order to report; "Causality is not a prerequisite for MedWatch reporting; suspicion that a medical product may be related to a serious event is sufficient reason for a health professional to submit a MedWatch report". Contributors are, however, warned that the system does not encourage 'a report on every adverse event' only serious events [255]. The aspirations for novel reporting systems, cited in previous paragraphs, must be balanced against the experience of many managers who are responsible for maintaining such systems. For example, the following quotation describes the results of an investigation into the processing of safety hazard reports and unsatisfactory condition reports (UCR). This was triggered by a letter that was sent by a Corporal in the Canadian National Defence Forces' Safety Digest publication in which he expressed frustration at the apparent failure of their reporting system:

> "The investigation revealed that, at least four years ago, Cpl Krygsveld's immediate supervisors failed to action or register four legitimate hazard reports that he submitted. Thus, the Unit Ground Safety Officer was not aware and no immediate action was taken. The fact that two supervisors tore up two of Cpl Krygsveld's hazard reports is unacceptable and clearly reflects a disturbing lack of safety consciousness." [138]

Although the investigation report argued that this is an isolated problem within the Canadian Airforce scheme, similar problems have affected many other reporting systems. For instance, Chapter 15 has described how the FDA have pioneered the development of Sentinel monitoring systems as a means of addressing the problems of under-reporting and reporting bias in accounts of medical device failure.

This following pages provide a critical analysis of recent initiatives to introduce incident reporting systems. In particular, it is argued that the proponents of many reporting schemes over-estimate the impact that they can have upon the operation of complex, safety-critical systems. One reason for this is that many recent proposals show a limited understanding of the tools and techniques that have supported reporting systems in other countries and other industries.

16.1 Human Problems

Many of the problems that affect the introduction and operation of incident reporting systems stem directly from the difficulty of encouraging individuals and teams to participate in the system. The scale of this problem is often under-estimated in the official announcements that are used to launch new reporting schemes. The publicity surrounding the launch of a new reporting system is often intended to help elicit participation. Unfortunately, such official announcements and press releases are insufficient to sustain most reporting systems.

16.1.1 Reporting Biases

The long-term elicitation of incident reports relies upon a number of factors. In particular, it can depend upon whether or not particular groups in the workforce perceive there to be any potential benefits from participating in the scheme. For example, nursing staff contributed about 90% of all of the reports that have been submitted in a local intensive care unit over the last decade. 621 reports were submitted by nurses compared with 77 reports by medical staff [119]. These figures must, however, be interpreted with great caution. It is important to consider the total number of staff who might contribute to such a system. The teams that contribute to this reporting scheme consist of three medical staff, one consultant, and up to eight nurses per shift. The larger number of reports contributed by nursing staff can also be explained in terms of the involvement in, or exposure to, the types of workplace incidents that were solicited under this particular scheme. Nursing staff had the most direct contact with the patients who remain the focus of the reporting system. They then have a proportionately greater opportunity to witness adverse events [119]. Such reporting biases have important consequences. The system may tell us a great deal about the execution of medical procedures. It may, however, tell us relatively little about more complex problems in the planning, coordination and administration of treatment within a department. Such factors are often overlooked by the proponents of reporting systems when they make strong claims about the quality and quantity of information that might be obtained about safety-critical systems.

Automated logging and tracking systems provide means of addressing the problems both of low contribution rates and of biased participation in a reporting system. The proponents of such systems often have an unfortunate way of advocating their introduction; “competent personnel love them, while incompetent personnel loathe them” [222]. Such assessments hide the difficulty and expense that is often involved in interpreting the data provided by automated, incident-monitoring systems. There is also a concern that any data will be used to punish rather than support staff performance through additional training. These concerns have acted as powerful barriers against the introduction of monitoring equipment onto UK trains. Recommendation 9 of the HMRI report into the accident at Watford South Junction advocated the use of automated systems to monitor driving technique. In 1999, however, less than 20% of trains carried this equipment [349]. More recently, the action plan to implement the recommendations of the Southall accident report included steps to extend both voluntary incident reporting systems and automated monitoring equipment [317]. This link between voluntary reporting systems and automated monitoring is instructive. The success of voluntary, confidential reporting systems in aviation, such as NASA’s ASRS, is often explained in terms of the pilot’s fear that any incident may have been observed and reported by their colleagues or by tracking equipment. The submission of a report provides them with a limited degree of legal protection and support in any subsequent investigation. It is less clear whether such a ‘no-blame’ approach might also be extended to other domains where automated monitoring and confidential reporting systems have been used together. For instance, some initial steps have been taken to use computer-based tools to automatically identify adverse occurrences in healthcare applications [400]. This often leads to ethical problems when failures have a long-term effect on the patient’s prognosis. How many times should a clinican be permitted to endanger their patients before we search for remedies other than re-training?

16.1.2 Blame

The US National Patient Safety Foundation (NPSF) recently argued that the names of individual health workers who file incident reports must not be released to the public or to licensing bodies. They maintained that ‘the culture of health care is rife with guilt, blame, and fear, and these are the greatest obstacles to effective reporting systems’ [584]. Similarly, the Australian Transportation Safety Bureau (ATSB) exploit a ‘no-blame’ systemic safety analysis that is intended to encourage a climate ‘in which people are prepared to report their errors’ [53]. In the UK, the National Occupational Safety and Health Committee has argued that ‘too often’ there is a tendency to blame the victim rather than search for the causes of adverse events [583]. ‘The most important thing’ is to establish why they were not prevented rather than focus narrowly on how they happened.

Again, however, many of these statements avoid the moral and practical problems that arise from the maintenance of ‘no blame’ reporting systems. The Cullen enquiry into the Ladbroke Grove accident acknowledges this complexity when considering the possible extension of no-blame, confidential reporting systems [194]. It was argued that people must be accountable for their actions. A no-blame approach can also, paradoxically, lead to a system in which workers are more ready to accept responsibility for a failure in order to conclude the investigation as rapidly as possible. Instead of no-blame reporting, the aim should be to develop an industry culture ‘in which information is communicated without fear of recrimination and blame is attached only where this is justified’ [194]. Cullen notes that the UK rail industry is some distance away from this ideal situation. Many large-scale systems, address employees’ fear of retribution and blame by developing elaborate legal safeguards to protect potential contributors [59]. Unfortunately, most reporting systems lack the funding and the necessary managerial support to provide this level of assurance. Even when they are present, there can be other workplace factors that dissuade employees from participating in a reporting system. In extreme cases, employees may even neglect medical treatment rather than expose themselves to workplace harassment:

> “FRA has become increasingly aware that many railroad employees fail to disclose their injuries to the railroad or fail to accept reportable treatment from a physician because they wish to avoid potential harassment from management or possible discipline that is sometimes associated with the reporting of such injuries. FRA is also aware that in some instances supervisory personnel and mid-level managers are urged to engage in practices which may undermine or circumvent the reporting of injuries and illnesses.” [233]

In the medical domain, a number of high-profile cases have acted as a powerful disincentive to the participation in incident reporting systems. For example, the High Court recently intervened to recommend the reinstatement of a surgeon who had expressing worries about the success rate of a colleague in his hospital. The trust initially refused to comply with the Court of Appeal’s finding [109]. This parallels the case of Stephen Bolsin who first uncovered an unusually high death rate among babies undergoing cardiac surgery at the Bristol Royal Infirmary [436]. He subsequently claimed that he was unable to continue working in the NHS as a result of his ‘whistle blowing’ and was forced to move to a hospital at Geelong, near Melbourne. These causes have resulted in the Public Interest Disclosure Act (1998), which allows whistle blowing staff who feel they have been victimised to take their employers to an industrial tribunal. There is no limit to the compensation that can be awarded and employees simply need an “honest and reasonable” suspicion that malpractice has occurred or is likely to occur. Such protection has, however, proven to be insufficient to persuade employees to contribute to many voluntary reporting systems. For example, the 2001 Royal College of Nursing congress explicitly backed a call for action to protect nurses who ‘speak out’. One of the delegates argued that whistle blowing was often seen as ‘grassing up’ or betraying colleagues. Theoretically, such additional protection should not be necessary under the 1998 Act. Some of these concerns can be explained by the informal pressures to conform to the norms of a particular working group. They can also be explained by the practical problems of preserving anonymity within small teams. Given the limited numbers of staff who perform particular tasks on particular shifts, potential contributors can often be identified through a simple process of elimination.

16.1.3 Analytical Bias

It is important not to underestimate the potential biases that influence the analysis of near misses and adverse occurrences. Over the past three years, we have conducted a series of interviews, surveys and observational studies of incident investigators and safety managers [749]. This work has helped to identify a range of influences that can affect the decision making processes that are intended to distinguish causal factors from the mass of other contextual information that is extracted from an initial report. At its most extreme, incident data can be used in a post hoc way to justify decisions that have already been made and positions that have already been adopted. For example, the proponents of Crew Resource Management training have used data from the US Aviation Safety Reporting System in this way [410]. Chapter 11 briefly introduced these sources of analytical bias.

For instance, *author bias* arises when individuals are reluctant to accept the findings of any causal analysis that they have not themselves been involved in. For instance, incidents at US highway-rail crossings can trigger investigations by federal organisations, such as the National Transportation Safety Board (NTSB) and the FRA. They can also result in state level enquiries. In some states, responsibility is divided between public agencies and the railroad operators. Elsewhere, responsibility is assigned to regulatory agencies such as the Public Utility Commission, Public Service Commission, or State Corporation Commission. In other states, investigations involve representatives of state, county, and city jurisdictions. Both state and local law enforcement agencies will also be involved if an incident involves the enforcement of traffic laws. Local government bodies are given responsibility for operational matters related to crossings through their ordinances. Each of these organisations can, and often do, hold different views about the causes of adverse events. The situation is slightly simpler for incident investigations in the UK. However, railway privatisation has created a situation in which conflict can arise between operating companies, Railtrack and HMRI. This is neatly encapsulated in Anthony Scrivener's recent article on Ladbroke Grove entitled 'Pass the signal - pass the blame' [732].

Confirmation bias arises when investigators attempt to ensure that any causal analysis supports hypotheses that exist before an incident occurs. In other words, the analysis is simply conducted to confirm their initial ideas. *Frequency bias* occurs when investigators become familiar with certain causal factors because they are observed most often. Any subsequent incident is, therefore, likely to be classified according to one of these common categories irrespective of whether an incident is actually caused by those factors [394]. There are many examples of these two forms of bias in the handling of SPAD reports prior to the Ladbroke Grove accident. Cullen estimates that approximately 85% of all such incidents were classified as the result of driver 'error' [194]. The frequency of such findings helped to reinforce this analysis as an acceptable outcome for any SPAD investigation; "I am led to conclude that the ready acceptance of blame by drivers, encouraged by the no blame culture, may have contributed to this poor analysis of root causes". The subsequent report argued that operating companies should review their incident investigation practices to ensure that there is no presumption that driver error is the sole or principal cause of SPADs.

Recognition bias arises when investigators have a limited vocabulary of causal factors. They actively attempt to make any incident 'fit' with one of those factors irrespective of the complexity of the circumstances that characterise the incident. These pressures can be illustrated by the response to initial reports of problems in the performance of cardiac surgery at Bristol Infirmary. The Society of Cardiothoracic Surgeons of Great Britain and Ireland discussed the reports of poor outcomes in 1989. Further information emerged during site visits in 1990. The sub-optimal results were attributed to the low volume of work because an increasing number of cases was widely believed to be associated with better outcomes. Adverse reports were, therefore, interpreted in a way that encouraged the generation of more work and that avoided questioning existing practices at lower volumes. The eventual enquiry argued that "the focus on throughput may with hindsight be thought to have distracted attention from further inquiry, as the Bristol results, with the exception of the figures for 1990, showed no real improvement" [436].

Political bias arises when a judgement or hypothesis from a high status member commands influence because others respect that status rather than the value of the judgement itself. This can be paraphrased as 'pressure from above'. *Sponsor bias* occurs when a causal analysis indirectly affects the prosperity or reputation of the organisation that an investigator manages or is responsible for. This can be paraphrased as 'pressure from below'. *Professional bias* arises when an investigator's colleagues favour particular outcomes from a causal analysis. The investigator may find themselves excluded from professional society if the causal analysis does not sustain particular professional practices. This can be paraphrased as 'pressure from beside'. The influence of these workplace issues can be difficult to assess. For example, the FRA Safety Board conducted an analysis of incidents from January 1990 to February 1999. This found that only 18 coded 'operator fell asleep' as a causal or contributing factor. The NTSB found these figures difficult to believe given the prevalence of such incidents in other modes of transportation [608]. Two NTSB investigations that had found fatigue as a causal factor were not coded in the FRA database as fatigue-related but as a failure to comply with signals. A number of influences might explain such different interpretations of the same incidents. For instance, the FRA plays a significant role in the promotion of the rail industry as well as in its regulation. The NTSB focuses more narrowly on the investigation of safety-related incidents. In consequences, the political, sponsor and professional influences that act on those organisations will be quite different.

This section has reviewed a broad range of 'human factors' problems that complicate the development and maintenance of incident reporting system. Some groups may choose to contribute to a scheme while others do not. Such participation patterns can be caused by a fear of retribution even in confidential no-blame systems. Analysts must, therefore, develop techniques to address the problems of under-reporting by eliciting contributions from potential participants. Alternatively, they can be forced to develop extrapolation techniques that can be used to make inferences about the nature of any potential incidents that might otherwise be reported by these groups. In either case, the problems of ensuring consistent participation across many different working groups can be exacerbated by the pressures that lead to analytical bias. For example, political bias can be exerted to ensure that a lack of reports from some groups is interpreted as positive evidence for a good safety record. Alternatively, high participation rates from certain groups can lead to forms of recognition bias that make analysts more likely to reach similar conclusions for incidents reported by different groups of workers.

16.2 Technical Problems

The interaction between particular participation patterns and the problems of analytical bias can frustrate attempts to obtain many of the potential benefits from incident reporting that were introduced in the opening sections of this Chapter. These problems would not be so serious if investigators and safety managers were equipped with an appropriate armoury of well-developed techniques. The methods and tools might then be applied to address the problems of under-participation. Individuals could then claim that any residual under-reporting had persisted 'in spite' of the most stringent efforts to elicit incident reports. Similarly, safety managers might claim that appropriate measures had been taken to combat various forms of analytical bias. Chapter 10 and 11 have described some of these techniques. Unfortunately, these techniques are not widely exploited. In contrast, most systems rely upon a range of ad hoc and 'in house' techniques to support both the elicitation and subsequent analysis of incident reports. This proprietary nature of these approaches can create barriers to information sharing. The use of different 'in house' methods also prevents comparisons being made between similar schemes in different countries or industries.

16.2.1 Poor Investigatory and Analytical Procedures

Previous chapters have identified a range of theoretical and practical issues that are often ignored during the development of small scale 'in house' analytical techniques. These issues, typically, have little importance during the initial stages of a reporting system. They can, however, become increasingly significant as the scope of any system expands to cover more potential contributors or as external regulatory intervention imposes increasing demands on those who are responsible for maintaining the reporting system.

We have identified two key theoretical ideas that must be considered when developing appropriate techniques for the analysis of adverse events: Mackie's Causal Fields and Hausman's Causal Asymmetries. Mackie argues that events result in effects that together form a 'causal field' [508]. For complex events, individual only observe a subset of these effects. Hausman's view of causal asymmetry builds on this argument [313]. If we know the cause we can predict the likely consequences, however, if we only know the consequences then it is far harder to unambiguously identify a single cause. An individual's interpretation of the cause of an incident, therefore, depends upon their observations of the effects and the relationship between those effects and a range of alternative possible causes. Additional complexity stems from the way in which most failures stem from several different factors that together form what Mackie terms a 'causal complex'. These theoretical ideas are reflected in the UK Health and Safety Executive's guidance on the incident and accident analysis that support railway safety cases:

> "There is much evidence that major accidents are seldom caused by the single direct action (or failure to act) by an individual. There may be many contributing factors that may not be geographically or managerially close to the accident or incident. There might also be environmental factors arising from or giving rise to physical or work-induced pressures. There is often evidence during an investigation that some of the contributory factors have been observed before in events that have been less serious. Accident and incident investigation procedures need to be sufficiently thorough and comprehensive to ensure that the deep-rooted underlying causes are clearly identified and that actions to rectify problems are carried through effectively." [350]

Unfortunately, many incident and accident techniques significantly under-estimate the complexity of causal analysis. Several existing approaches attempt to identify a single 'root cause'. Other techniques, fail to consider the range of alternative causes that can account for the same observed effects. This creates problems during subsequent enquiries and litigation when it can be shown that investigators failed to consider other plausible accounts. Such caveats can be levied at some of more advanced analytical tools, including WBA and Tripod.

There are a number of reasons why reporting systems fail to adopt existing investigation and analysis techniques. This book is, in part, intended to address the lack of reference material in this area. There are other problems. For instance, many small scale systems lack the resources that are necessary to hire or train existing staff in some of the more complex techniques. This creates particular problems when safety managers and investigators consider human factors issues. There is a tendency to blame incidents on inadequate attention or on poor staff performance. Such findings obscure or neglect the 'performance shaping factors' that contribute to human failure. This can be illustrated by Busse and Wright's analysis of incidents reported in an intensive care unit. The clinicians and nursing staff who were responsible for the system argued that a number of incidents stemmed from inattention and 'thoughtlessness'. This often led to recommendations that focussed on reminders, including numerous posters that describe recommended procedures [121]. The same events were then analysed by a human factors expert who argued that such reminders could only provide short-term protection against certain classes of adverse events. Their effectiveness declines rapidly over time. In contrast, the application of incident investigation techniques derived from the Tripod method, introduced in Chapter 11, revealed that many of the incidents of 'thoughtlessness' could also be interpreted as the result of 'work arounds' to support poorly designed or faulty equipment.

Larger-scale reporting systems can avoid some of these problems by ensuring that their staff are trained in appropriate analytical techniques. Unfortunately, there is little agreement about which approaches might support the causal analysis of incidents in many industries [453, 194]. It is instructive to note that even the GAIN initiative, which many regard as the most advanced attempt to create industry-wider reporting standards in aviation, has still to agree on a core set of analytical techniques. This lack of consensus has important consequences. It can undermine confidence in the findings of any investigation, especially when there are misgivings about the intent or purpose of any enquiry.

16.2.2 Inadequate Risk Assessments

It can be argued that safety managers and investigators are justified in their decision to reject many existing analytical tools in favour on 'in house' solutions. Previous chapters have argued that very few of the existing techniques can be integrated directly into the design and development of future systems. In particular, they are very poorly integrated with risk assessment. This lack of integration can have unfortunate consequences. For instance, I recently witnessed a design team deriving rough reliability estimates for the same components that their colleagues had already been studying using automated monitoring systems [423]. Incident reporting systems can provide evidence about the consequences of a potential failure and approximations for the likelihood of particular hazards. Such information can help to increase the accuracy of risk assessments which can be notoriously inaccurate [249].

The lack of integration between risk assessment and incident reporting not only affects the pro-active use of failure information to support future development. It can also prevent the effective allocation of resources within a reporting system. If investigators do not assess the risks associated with the recurrence of a previous incidents then it can be difficult to justify why one failure deserves closer investigation than another [416].

One of the main conclusions from this book is, therefore, that more support must be provided to support the two-way flow of information between risk assessment and incident investigation. The products of risk analysis can be used to gude the allocation of investigatory resources. The products of incident reporting can be used to inform estimates about the consequence and likelihood of future failures. A number of problems complicate the use of incident reporting data to guide the application of risk assessment techniques. Most risk analysis centres on the frequency and consequence of an event. It is, however, often unwise to assume that any recurrence of a near-miss incident will have the same outcome. Many reporting systems therefore assume that any recurrence will have the 'worst plausible outcome' [423]. This creates problems because different investigators can have very different opinions about what is, and what is not, a plausible outcome from any future failure. This is most apparent in the differences that can arise between the risk assessments that are produced by the safety managers in operating companies and those of regulatory organisations. The US Department of Energy will issue a Preliminary Notice of Violation as a way of warning management that they have under-estimated the risks associated with any recurrence of a particular incident. For example, a series of 'unplanned worker contaminations' in a national laboratory during 1999 led to exposures that were well within specified limits. However, the Department concluded that 'the lack of adherence to radiological work controls and the amount of radioactive material potentially available for uptake in the body' created the potential for more serious incidents in the future. The subsequent investigation argued that, in contrast, 'laboratory management was reluctant to acknowledge the serious nature of the concerns and treated them as a series of individual personnel errors' [206].

As mentioned, few analytical techniques provide explicit support for the use of risk assessments to drive the allocation of finite investigatory and development resources. Even if risk assessments are integrated into other investigatory techniques, there is no guarantee that investigators will respond in an appropriate manner:

> "During the almost five years preceding the Ladbroke Grove accident, there had been at least three occasions when some form of risk assessment analysis on the signaling in the Ladbroke Grove area has been suggested or proposed. The requests were: the Head of Technical Division's letter of 11 November 1996 which requested a layout risk assessment of the re-signaling (paragraph 43); the Field Inspector's letter of 16 March 1998 to Railtrack (paragraph 64); and the Railtrack Formal Inquiry of 1 July 1998 (paragraph 66). In addition there was an earlier request for details of measures taken to reduce the level of SPADs in the area around SN109 recorded in the Head of Technical Division's letter of 1st March 1995 (paragraph 39). None of these requests appears to have been pursued effectively by HMRI." [351]

Such comments illustrate another of the numerous paradoxes that arise in incident reporting. It is easier to identify situations in which risk assessment has failed to prevent a recurrence than situations in which it has successful mitigated the risks of future failure. For instance, safety managers might use a risk assessment to justify intervention to mitigate the consequences and likelihood of a particular failure. If their intervention has been successful then the number of similar incidents may fall and the outcomes of these events will be less 'severe'. However, such situations are indistinguishable from those in which the manager introduces unnecessary measures to address a risk that was lower than they had anticipated.

16.2.3 Causation and the Problems of Counter-Factual Reasoning

Investigators must adopt a consistent approach to the causal analysis of adverse events. Confidence can be compromised if different causes are identified for apparently similar events. Unfortunately, the proponents of many reporting systems underestimate the difficulty of causal analysis. The previous section outlined some of the theoretical problems that complicate this task. Counterfactual reasoning provides the main analytical technique for improving the consistency of causal analysis in incident and accident investigations [248, 469]. Chapters 10 and 11 have described how this technique has been integrated into a wide range of methods including Events and Causal Factor analysis as well as WBA. Counterfactual reasoning takes the general form that 'if a causal factor had not occurred then the incident also would not have taken place' [491]. If an incident would still have taken place whether or not a event had occurred then it cannot be thought of as causal factor. This style of argument is illustrated by an NTSB marine incident; "...had the main switchboard been subjected to thorough and timely inspections as part of an effective preventive maintenance program, any faulty connections or conductive objects would have likely been identified and corrected, and the fire might have been avoided." [618]. The same style of reasoning can be used beyond the immediate 'causes' of an incident to look at the actions that might have mitigated the consequences of the failure. For example, the report argued that "a firefighting team that was trained in the techniques of combating an electrical fire should have led the response to the fire in the control room...such a team probably would have extinguished the fire more quickly and with minimum risk". The author is using a counterfactual argument because a trained firefighting team was not available to combat the initial incident. The NTSB investigator also deploys counterfactual argumentation to eliminate potential causal factors. The report argues that 'even without a fuse, a transient voltage spike of sufficient magnitude to create an arc that could jump the gap probably could not have been created'. This is a counterfactual argument because there was evidence to suggest that a form of fuse had been present in the system before the incident. Hence with this additional safeguard, we can discount the cause of the fire being a transient voltage spike.

Counterfactual reasoning is both complex and error prone. For example, how sure can we be that an incident would not have occurred if a causal factor had not been present? Causal asymmetries suggest that many different causal complexes will have the same outcome. For instance, there are no guarantees in the previous incident that the inspections would have found a particular faulty connection. Previous incidents have shown that inattention and fatigue often compromise such safeguards. Chapter 11 has argued that the strengths and weaknesses of counterfactual reasoning remain an area for future research. Byrne and her colleagues have, however, conducted a number of preliminary studies [123, 124]. This work argues that deductions from counterfactual conditionals differ systematically from factual conditionals and that, by extension, deductions from counterfactual disjunctions differ systematically from factual disjunctions. This is best explained by an example. If we argue that "...had the main switchboard been subjected to thorough and timely inspections as part of an effective preventive maintenance program, any faulty connections or conductive objects would have likely been identified and corrected, and the fire might have been avoided" then readers will infer that the inspections had not taken place. This counterfactual style of argument can have such a persuasive effect that readers overlook contradictory evidence elsewhere in a report [426]. There are more complex examples of the inferences that readers draw from counterfactual arguments. The statement that *the fire was caused by a faulty connection within the main switchboard that initiated an arc fault or by a conductive object falling onto the switchboard bus bars* is a factual disjunction. Byrne argues that such sentences encourage the reader to think about these possible events and decide which is the most likely. There is an implication that at least one of them took place. The statement that *had the switchboard been covered by an effective preventive maintenance program or a thorough inspection by the Alaska Marine Highway System then the presence of faulty connections would have been identified* is a counterfactual disjunction. Chapter 10 has shown that this use of the subjunctive mood communicates a presumption that neither of these events actually occurred.

This theoretical work has pragmatic implications for incident investigation. If factual disjunctions are used then care must be taken to ensure that one of the disjuncts has occurred. If counterfactual disjunctions are used then readers may assume that neither disjunct has occurred. The distinction between counterfactual and factual disjunctions forms part of a wider concern to ensure that analytical biases are not hidden through the inappropriate use of language in incident reports. For example, rhetorical devices known as tropes can be used to increase the impact and effectiveness of everyday prose. Chapter 13 has briefly introduced the way in which tropes can be used to achieve particular effects on the readers of an incident report. For instance, *anaphora* uses repetition at the beginning of successive phrases, clauses or sentences. It can create an impression of climax in which the repetition leads to a particularly important insight or conclusion.

> "Both patients had implanted pacemakers, and both had experienced unintended maximum pacing rates up to 120 beats per minute. Medical intervention was needed to turn off the minute ventilation sensor in each pacemaker. When the sensors were turned off, the patients' heart rates returned to normal." [273]

This example illustrates the successive use of 'both' to emphasise the link between events happening to the patients. The investigator uses this repetition to draw the reader's attention to relationships between the consequences of a single cause for both patients. In this case, a clinical device such as a cardiac monitor or mechanical ventilators, was assumed to have generated a weak electrical signal that was sufficient to interfere with the ventilation sensors on the patient's devices. This, in turn, resulted in the incorrect measurement of thoracic impedance and ultimately in pacemaker rate increases. It is important to emphasise that such techniques are not of themselves either 'good' or 'bad'. Rhetorical devices can be used to convince us of well-justified conclusions or to support half-baked theories. It is important, however, to be sensitive to the effects that such techniques might have on the readers of an incident report. For instance, the previous citation can be interpreted to provide readers with a clear summary of the evidence that supports the investigators' conclusions. It can also be interpreted in a more negative light. For example, further investigation might establish independent causes for the effects observed in both patients. The rhetorical device creating a link between each individual might dissuade investigators from conducting such additional investigations.

Antithesis uses juxtaposition to contrasts two ideas or concepts. This technique is often used to contrast some form of normative or correct behaviour with the events that are presumed to have caused an accident. This can be illustrated from the following analysis of an incident reported to the Canadian Defence Forces:

> "...instead of braking gently, the driver's foot accidentally hit the accelerator. The vehicle jumped forward out of control, veered to the right, sped over the ditch, and crashed into the front wall of a 7-unit multi-family dwelling..." [147]

This technique is important because readers may make a number of additional inferences based upon such constructions. In this context, it is tempting to infer that the resulting collision would not have happened if the accelerator had not been pressed. The rhetorical construct diverts attention away from alternative hypotheses. The car may have been travelling too fast for any braking manuever to have prevented the eventual incident.

Most investigators and safety managers are unaware that they exploit such rhetorical devices. They draft prose to support their arguments. They may inadvertently stress conclusions that are not well supported by the available evidence. They may also cast doubt on other findings that contradict their version of events. Unfortunately, this inadvertent use of rhetorical devices is often exposed at litigation. In particular, it is often possible to show that particular linguistic constructs reflect the unsupported assumptions of investigators. In the previous example, it would be necessary to demonstrate that the accidental use of the accelerator was the cause of the incident and not the failure to brake well before the accelerator was applied. The key point here is that the problems of bias and interpretation not only affect the causal analysis of incident reports, they also complicate the way in which adverse events are documented and presented by investigators. If these influences are not considered then there is a danger that alternative explanations will be prematurely discounted and potential lessons lost.

16.2.4 Classification Problems

It can be difficult to detect patterns of failure amongst the natural language accounts of adverse events that are produced by many reporting systems. The volume of prose produced in national and international systems can make it difficult for any individual to keep track of common causes or consequences across many incidents. In consequence, many reporting systems use keyword-based summaries. Analysts represent the causes of an incident or near-miss by an enumeration of terms drawn from an agreed glossary. This approach can also provide a concise representation of a range of other contextual information, including mitigating factors and the potential consequence of an adverse event. The use of keyword summaries helps to reduce the interpretation problems that stem from tropes in natural language accounts. This approach strips out the rhetorical techniques that emphasise particular interpretations through the use of anaphora, antithesis etc. There are further benefits. For instance, the use of an agreed taxonomy can help to ensure that different organisations all consider a consistent set of terms when describing adverse events. Chapter 14 has also shown how classification schemes can be based upon the data models that support relational databases. Not only do these terms provide a vocabulary for describing individual incidents, they also provide the keywords that can be used to form the queries used to extract information about previous events. This use of classification schemes also supports the compilation of statistical data. Safety managers and regulators can provide information about the frequency of incidents that are attributed to each of the causes included in the taxonomy.

A number of practical problems complicate the use of such taxonomies to support the indexing and retrieval of incident reports. In particular, it can be difficult to establish reliable procedures for the codification of each adverse event. Each incident can be codified locally, within the group or organisation in which an incident occurred, or by a central unit who are responsible for codifying a large number of events sent in by different participants. If the codification of incidents is performed centrally then it is important that staff understand enough about the context in which an incident occurs for them to ensure that the correct codes are assigned. Alternatively, if incidents are to be codified at a local level then it can be difficult to ensure that safety managers assign the same codes to similar incidents in different locations [417]. The FDA illustrate some of t he problems in incident classification through a 'real' case study in which a violent patient in a wheelchair was suffocated through the use of a vest restraint that was too small. The risk manager, JC, proceeded as follows:

> "She finds the list of event terms, which was detached from the rest of the coding manual... She muses: 'Mr. Dunbar had OBS which isn't listed in these codes; he had an amputation which is listed; he had diabetes which isn't listed; and he had hypertension which is listed'. JC promptly enters 1702 (amputation) and 1908 (hypertension) in the patient codes. She then finds the list for Device-Related Terms... She reviews the terms, decides there was nothing wrong with the wheelchair or the vest restraint, and leaves the device code area blank." [275]

The success of any classification system, therefore, depends upon the procedures that are used to identify appropriate terms. The resulting classification of 1702 (amputation) and 1908 (hypertension) provides few insights into the nature of the incident. A range of techniques can be used to identify potential mis-classifications. Many of these rely upon comparing the results from any classification with those obtained from more reliable sources. For example, the frequency of particular terms in an incident classification can be compared with those from Sentinel schemes. The additional resources and training provided to Sentinel systems should ensure that their classifications provide a more accurate reflection of adverse events than those submitted by other reporting systems. Unfortunately, it can be difficult to judge whether any differences are due to misclassifications or to underlying differences in the nature of events that are reported to different units within the same industry. Alternatively, analysts can compare the results of an independent reclassification of previous incidents with those that were originally returned from a reporting system. Such comparisons again rely upon an appropriate sampling technique to ensure that this approach detects 'genuine' differences in any subsequent reclassification. If the sample focuses on particular classes of adverse event then this approach may fail to uncover wider problems of misclassification in the incidents that wre not selected as part of the sample.

A range of further problems complicate the application of taxonomies. Many incidents involve 'wicked' failures that cannot easily be described by a number of discrete terms [468]. For instance, computer-related failures often stem from a combination of requirements and design errors. System components fail because some necessary tasks are not identified and because others are identified but had not fully developed. For complex systems it can also be difficult to distinguish between a requirements failure and a design failure. This is especially difficult if detailed design and requirements documents are distributed across the sub-contractors that are responsible for implementing component functionality.

Further problems arise from the difficulty of classifying human behaviour. For instance, a recent HMRI report identified the difficulty of distinguishing between 'misjudgement' and 'disregard' [349]. The allocation of these different terms has a profound impact on the consequences of any investigation. 'Misjudgement' implies that the operator may have behaved in a reasonable manner even if they ultimately failed to safeguard the system. 'Disregard' suggests a more willful neglect of necessary precautions. The HMRI report cites the example of a train driver who appeared to make every effort to brake at a signal in poor weather conditions, yet the incident was categorised as 'disregard' rather than 'misjudgement'. The allocation of such terms involves a level of analysis and discretion that goes beyond the FDA's taxonomy, which focuses on observable features such as the role of hypertension in an incident.

In many industries, it can be difficult to ensure agreement over the components of incident taxonomies. In particular, there is a trade-off between the coverage of a taxonomy and the reliability of any analysis. In general, analysts are less likely to achieve a consistent classification if more terms are introduced into a scheme. The development of an appropriate taxonomy is further complicated by the need to respond to changes in the types of incident that are reported to a scheme. This creates particular problems if analysts are forced to go back and manually reclassify hundreds or thousands of records to reflect new distinctions and definitions of the components of any taxonomy. Some reporting schemes now hold more than 500,000 reports [59]. If previous records are not updated to reflect the new classification system then safety managers may fail to discern that recent incidents form part of a wider pattern, which is obscured by weaknesses in the previous classification scheme. This problem is particularly acute when taxonomies are extended to describe human behaviour. The field of human factors research has changed rapidly over the last decade with an increasing focus on group interaction. However, few of these changes have been reflected in incident reporting systems because of the costs associated with manually analysing and re-classifying existing records.

As mentioned, many classification systems are derived from or inform the development of databases. These systems support the retrieval of individual incident reports using queries that are constructed in terms of the components of incident taxonomies. For instance, analysts can use the FDA's MAUDE system to retrieve information about all incidents that were tagged with the device codes, mentioned in previous paragraphs. Unfortunately, the theoretical under-pinnings of these systems are often poorly understood by the people who use them. Safety managers, therefore, often rely upon queries that are pre-programmed by system administrators. Unfortunately, it can be difficult to ensure that safety managers clearly communicate their information requirements to technical support staff [413]. In consequence, safety managers often do not receive the information that they think they have requested when they issue a query. It can also be difficult for safety managers to formulate more ad hoc, exploratory queries because they lack necessary technical knowledge about relational database technology [413].

Chapter 14 introduced alternative technologies that avoid some of the limitations associated with the use of relational databases in incident reporting schemes. Web-based techniques can also be used to automate the indexing of incident reports in response to changes in a classification scheme. This avoids the overheads associated with the manual reclassification of many thousands of previous records. Similarly, probabilistic information retrieval systems enable users to search for information without the need to form complex queries [413]. Information requests can be expressed in the vernacular. The retrieval system analyses compares attributes of the query, such as the frequency of key terms, to automatically identify potential matches within a collection of incident reports. For example, it is relatively rare to find the term 'explosion' in medical incident reports. The use of this term in a query can help retrieval systems to identify a relatively small number of potential matches. In contrast, less attention might be paid to the use of the term 'patient', which is likely to appear in many of the records within the system. Probabilistic information retrieval systems do not, however, provide a panacea. For example, it can be difficult to ensure that particular queries yield appropriate levels of precision and recall. A system exhibits poor precision if it returns many incidents that the user does not believe are related to their query. The user must then manually filter the large number of incidents that the system considers to be a match. Conversely, poor recall occurs when a system fails to return an incident that the user believes is related to their query. Systems that provide good recall are often imprecise. Conversely, systems that offer high degrees of precision will often exclude incidents that ought to have been returned as a potential match. A number of existing research projects have, therefore, begun to look for alternative computational techniques including conversational case based reasoning [413]. The intention is to avoid the static data models that limit the application of relational databases but also to help users interactively address the problems created by poor precision and recall. Conversational case based reasoning techniques enable users to filter the presentation of incident information through the iterative refinement of natural language queries.

16.3 Managerial Problems

Previous sections have focussed on problems that can prevent reporting systems from yielding the benefits that are often claimed for them. These range from the difficulty of eliciting sufficient reports through the problems of ensuring consistent analysis through to the computational challenges that complicate the storage and retrieval of large-scale incident collections. In contrast, this section focuses on the problems of managing incident reporting systems. It is important not to underestimate these problems because they are often far more significant than the technical issues that have been summarised in previous paragraphs. If insufficient resources, including time, finance and expertise, are allocated to a reporting system then there is little prospect that it will yield significant safety improvements. If management structures are not established to enable the dissemination, implementation and monitoring of safety recommendations then any insights are unlikely to be acted upon. If higher levels of management, regulators and political interests have unrealistic expectations about the likely benefits of a reporting system then it is unlikely to satisfy any subsequent validation actions.

16.3.1 Unrealistic Expectations

Chapter 1 has presented statistics that indicate the high frequency of adverse events and near-miss incidents in many safety-critical industries. It has been estimated that between 4% and 17% of patients in acute hospitals suffer from iatrogenic injury [849]. Observational studies have found that 45% of patients experienced some medical mismanagement and 17% suffered events that led to a longer hospital stay [28]. The US Aviation Safety Reporting System averages approximately 600 reports and UK railway's Safety Management Information System receives over 1,700 reports each week. It can, however, be difficult to interpret these statistics. The 4-17% figure for iatrogenic injury is based on extrapolations that may not be confirmed by new regional reporting procedures in the UK and the US. Similarly, the number of submissions to the Safety Management Information System is partly due by the way it helps record reports that meet regulatory requirements under the UK's RIDDOR regulations. In spite of these caveats, many of the proponents of incident reporting point to the success of schemes such as the ASRS to justify the development of new systems. They often underemphasise the potential problems that can arise when data is collected unnecessarily or that can occur when finite resources are swamped by a flood of 'low criticality' reports. For instance, since the Ladbroke Grove accident, there has been a requirement to conduct a formal inspection after every report of a SPAD. This has resulted in over 200 formal investigations by Her Majesty's Railways Inspectorate (HMRI). It is likely that there have also been a far larger number of near-misses. Drivers frequently avoid passing the signal by rectifying a potential problem 'at the last minute'. It has been argued that confidential, voluntary reporting systems might be used to elicit information about these events that would otherwise go unrecorded [194]. In particular, they might provide insights both about those measures that helped the driver to detect the potential danger and about those factors that might have turned a near miss into a more serious incident.

Some individuals and organisations have looked beyond the particular safety lessons that might be learned from a reporting system. In addition to identifying successful defences and potential vulnerabilities, they have argued that such schemes will also reduce costs by avoiding the negative consequences of previous failures. For instance, the NHS have promoted the development of voluntary reporting systems as one of several measures that are intended to achieve a number of ambitious objectives:

> "...the Department of Health should establish groups to work urgently to achieve four specific aims: by 2001, reduce to zero the number of patients dying or being paralysed by maladministered spinal injections (at least 13 such cases have occurred in the last 15 years); by 2005, reduce by 25% the number of instances of negligent harm in the field of obstetrics and gynaecology which result in litigation (currently these account for over 50% of the annual NHS litigation bill); by 2005, reduce by 40% the number of serious errors in the use of prescribed drugs (currently these account for 20% of all clinical negligence litigation); by 2005, reduce to zero the number of suicides by mental health inpatients as a result of hanging from non-collapsible bed or shower curtain rails on wards (currently hanging from these structures is the commonest method of suicide on mental health inpatient wards)." [633]

Such high expectations can contrasted with the prosaic problems that limit the utility of incident reporting systems. There are further problems. Chapter 15 has described the difficulty of monoring whether or not a reporting system is meeting the ambitious objectives that often set for them. Increases in the reporting frequency can be due to increased participation in a scheme or from a sudden rise in adverse events. Independent audits can be used to distinguish between these two interpretations, however, limited resources prevent their use as a persistent monitoring device. There are further problems; a reduction in the accident frequency cannot always be used as an indicator. In many industries, high consequence adverse events are so rare that any reduction from the beneficial effects of incident reporting cannot be distinguished from random changes even over relatively long periods of time. A number of more sophisticated statistical models have recently been proposed to address these concerns. For instance, UK RailwaySafety's Precursor Indicator Model measures short-term changes in the frequency of events that have led to previous accidents. The precursors that form the focus of this monitoring technique are updated over a longer time period to reflect any changes in the causes of those accidents that do occur. Such approaches recognise the technical complexities that arise in attempting to extract data to support the management of incident reporting systems. One of the aims behind this book is to undermine some of the more simplistic and extreme claims that are made about the benefits of incident reporting. A further intention has been to provide safety managers with information about tools, such as the Precursor Indicator Model described in Chapter 15, that can be used to avoid some of the pitfalls that have affected previous reporting systems.

16.3.2 Reliance on Reminders and Quick Fixes

Risk monitoring tools, such as RailwaySafety's PIM techniques, are not an end in themselves. They provide information that should be used to inform future decision making. Information about previous failures must guide the interventions that reduce the likelihood or mitigate the consequences of future incidents. Previous chapters have cited many success stories where incident reports have triggered a prompt and effective response to adverse events. These have ranged from new forms of glazing in the viewports of military bunkers described in Chapter 12 [144] through to changes in the ventilation systems in the laundry facilities on cruise ships mentioned in Chapter 13 [606]. It is important to recognise, however, that the development of a reporting system does not always guarantee that recommendations will be acted upon. For example, Chapter 9 recounted the thirty year campaign by the NTSB to encourage the wider use of Excess Flow Valves in consumer gas lines [588]. Similarly, the concerns voiced by the ASRS about the consequences of tight turnaround times in US airports have had some limited effects [410]. It is unrealistic to expect that every recommendation proposed by a reporting system should be implemented. Economic and commercial pressures have been cited as reasons for the delay in implementing the two recommendations mentioned above.

There are further generic reasons why reporting systems can fail to yield the safety benefits that some have predicted. For instance, these schemes often yield few surprises. They help to reinforce existing safety concerns that many managers will already be very familiar with. For instance, Sexton, Thomas and Helmreich's study of medical incident reporting found that the most common recommendation for improving patient safety in intensive care was to acquire more staff to handle the workload. The most common recommendation from incident reporting in operating theatres was to improve communication [735]. The lack of surprise should be unsurprising. Organisations that establish reporting systems, typically, already have a good idea of the safety issues that affect their working practices. Van Vuuren's research, cited in Chapter 11, provides further examples. One study focussed on 19 incidents that were reported over approximately one month to an Accident and Emergency Department [844]. His analysis yielded a total of 93 potential causes of which 45% related to organisational issues while 41% were classified as 'direct' human causes. The organisational causes included the need to secure external services. In particular, incidents were often triggered or exacerbated by the need to secure beds for the patients in the Department. They also included a lack of senior staff during peak periods. Direct human causes included problems that new Senior House Officers experienced in interpreting X rays. They also stemmed from a culture of learning from mistakes and a reluctance to contact senior staff. Many NHS safety managers are already very aware of these issues. Such concerns do not, however, secure the resources and organisational support that is necessary to implement specific improvements.

It is difficult to underestimate the impact that resource issues can have upon the benefits of any reporting system. Some schemes have proven to be very successful in triggering large scale investments in safety measures. Insurance and safety classification companies, including Lloyds Register and DNV, have played an important role in motivating upper levels of management to invest in the recommendations that are derived from reporting systems. In contrast, other more local schemes have often been set the target of becoming 'self funding'. In other words, any safety investments must be paid for by corresponding process improvements that are elicited as 'lessons learned' through the same reporting mechanisms [417]. These constraints can often lead to a form of resource starvation. Managers are forced to make recommendations that they know will never be funded or to focus their attention on 'low cost' and 'no cost' solutions. This partly helps to explain the dominance of the *perfectability approach* to risk management, mentioned in Chapter 12. Rather than address the organisational, technical and environmental causes of adverse events, staff are urged to 'try better', 'be more aware' or 'follow established procedures'. Busse and Wright's study of reporting in intensive care found 82 'Remind Staff?' statements in a total of 111 recommendations over a 15 month period [119]. 29 other recommendations focussed on revised procedures and protocols (e.g. 'produce guidelines for care of arterial lines - particularly for femoral artery lines post coiling'), or were equipment related (e.g. 'Obtain spare helium cylinder for aortic pump to be kept in ICU'). A reliance on reminders can also be seen in larger scale, voluntary reporting systems. For example, Chapter 2 described how the aviation industry's CHIRP system often relies upon the perfectability approach. Pilots have been urged to check that they have entered the 'correct' data into navigation systems 'then, and only then, should the Execute function button be pressed' [176]. Similarly, they have been urged to complete pre-flight visual inspections of all flight surfaces or in cases where this is not possible, such as a high wing high tail configuration, to ensure 'a sound knowledge' of local de-icing processes in the prevailing weather conditions, 'If there is any doubt as to whether the aircraft is clean, a take-off should not be attempted' [179]. Similarly, the ASRS has encouraged pilots to mitigate the problems of noise cancelling headphones by having them 'half-on' and 'half-off' during take-off and landing [62]. These reminder statements can be interpreted in two different ways. At one level, they provide an important source of practical information for operators. This is reinforced by a recent contributor to CHIRP:

> "It is timely that we remind ourselves of the health and safety hazards that may exist on the aircraft. It is also timely that we remind ourselves that we are individually responsible for our own health and welfare in situations that we know are hazardous. [178]

These reminders can, however, be interpreted in a different way. It is possible to argue that far from improving safety, they illustrate potential weaknesses in the defences that protect safety-critical systems. Human factors research points to the dangers of any reliance on reminders. Unless people are continually reminded then they are likely to forget the importance of safety precautions over time [367]. Any reminder should be seen as a prompt for more concerted action to address underlying technical and organisational issues. In this view, reminders are short-term fixes to deeper safety problems ranging from the design of human-machine interfaces to navigation systems through through to the intergration of noise cancelling headphones into cockpit auditory warnings. There is evidence that the managers of both the CHIRP and ASRS systems have responded in this way, they have initiated and contributed to more sustained safety initiatives following particular incident reports. It is also instructive, however, to examine those incidents that typically yield reminders. As we have seen, in the aviation domain these often concern cockpit design but can also relate to commercial pressures to meet particular ATC slots. In the medical domain, reminders are often used to cope with the lack of beds or key staff or with specific medical devices acquired by other units in the hospital. These reminders are intended to address problems that are, typically, perceived to lie outside the scope of the reporting system. Staff must perform better to cope with issues that cannot be addressed by more direct means, including changes in acquisition policy and commercial practices. The presence of such reminders, therefore, provides insights into the perceived limitations of an incident reporting system.

16.3.3 Flaws in the Systemic View of Failure

The reliance on reminders by ‘perfectionist’ approaches to incident reporting often stems from an undue emphasis being placed upon the direct human causes of adverse events. Operators are blamed as the main cause of an adverse event, hence they should be exhorted to ‘try better, next time’. Correspondingly less attention is placed on the organisational, technical and environmental circumstances the contributed to an incident. The managers of many schemes have become sensitive to these criticisms. For instance, Chapter 13 quotes the Swedish Maritime Agency’s statement that ‘it is to be underlined that it is not the purpose of the investigation to establish or apportion blame or liability’ [768]. Rather than focusing on individual human errors, the purpose of their reporting system is to provide ‘a complete picture’ of adverse events. Similarly, a recent report into SPADs on UK railways stressed that ‘no driver sets out to have a SPAD, but all humans are prone to unintentional error on occasions’. It went on to argue that in order to reduce SPADs it was necessary to understand the factors that contribute to these events and that ‘... there is a growing body of evidence that features in the design and configuration of the signalling system can significantly increase the risk of driver error’ [356]. This view has also been embodied in HMRI guidance on the preparation of safety cases, this rejects the identification of operator error as a root cause of incidents on UK railways:

> “In these criteria the term ‘root causes’ includes consideration of management’ s real and perceived messages to workers, environmental and human factors, as we ll as plant failures and inadequate procedures. Human errors arising from poor operating conditions, procedures, management expectations or plant design are not root causes; the predisposing factors are.” [350]

In the domain of medicine, the UK NHS’ ‘Organisation with a Memory’ argued that although human error can precipitate an incident, there are usually deeper ‘systemic factors’ that created the context in which an adverse event was likely to occur [633]. In contrast to the perfective approach, this systemic approach:

> “... takes a holistic stance on the issues of failure. It recognises that many of the problems facing organisations are complex, ill-defined and result from the interaction of a number of factors. This approach starts from the premise that humans are fallible and that errors are inevitable, even in the best run organisations. Errors are seen as being shaped and provoked by upstream systemic factors, which include the organisation’s strategy, its culture and the approach of management towards risk and uncertainty. The

> associated counter-measures are based on the assumption that while we cannot change the human condition we can change the conditions under which people work so as to make them less error-provoking. When an adverse event occurs, the important issue is not who made the error but how and why did the defences fail and what factors helped to create the conditions in which the errors occurred. The system approach recognises the importance of resilience within organisations and also recognises the process of learning as enhancing such resilience." [633]

The claim that accidents occur in the 'best run organisations' echoes Perrow's argument that companies will still have accidents no matter how hard they try [675]. His work on 'normal accidents' provides theoretical under-pinning for the systemic view of failure [675] Accidents occur in unexpected ways because operators, managers and regulators cannot hope to control the many different hazards that are created by the complexity and coupling of high-technology systems. It is easy to misinterpret many of Perrow's arguments especially about the nature of politics and control in high-reliability organisations. It seems clear, however, that he views adverse events as part of the price to be paid for technological innovation. Accidents and incidents also provide 'warning signals' about the trade-offs between safety and production that characterise many industries. The previous chapters of this book provide numerous examples to confirm Perrow's view. Railtrack have made considerable efforts to introduce innovative incident reporting systems. They have also been involved in a succession of major accidents. Perrow argues that if we cannot prevent potential catastrophic failures in high-technology systems then they should be 'abandoned, drastically scaled back or drastically redesigned'. However, many of these rail accidents occurred while design changes are being introduced across the network. The introduction of the Train Protection and Warning System, mentioned in Chapter 3, is scheduled for completion in January 2004. Unfortunately, it is difficult to abandon or drastically scale back a national rale system in the meantime. The reduction in network traffic following the Hatfield crash played a major role in the UK governments decision to force the infrastructure company into receivership.

At the heart of normal accident theory is the argument that previous technologies supported linear and loosely coupled systems. It is possible to anticipate and counter the hazards that are raised by dams and canals. Accidents are 'foreseeable and avoidable' because if one component fails then there is time to react and mitigate the consequences of adverse events. In contrast, modern production and transportation systems exhibit non-linear behaviours. They exploit techniques such as 'just in time' production that provide efficiency savings but that also create potential vulnerability. Individual failures propagate throughout and beyond system components in ways that would not have been possible in loosely coupled systems. Again, previous chapters have supported these arguments. For instance, Chapters3 and 7 describe how a decision to increase the traffic on a Brunswick Line rail system led to the placement of a signal before a station so that drivers had sufficient time to break before a hazard on the other side of the station. This contributed to the subsequent accident because investigators argued that drivers were less likely to recollect the previous signal after they had performed a station stop [596]. This incident illustrates many of Perrow's arguments. The decision to increase line capacity implied higher operating speeds and longer stopping distances. This consequent operating changes introduced greater complexity for drivers and signallers. The increased operating speeds reduced the opportunity to correct any driver failure to recollect the signal. High-line capacity also made it more likely that any such failure would have catastrophic consequences; it was more likely that the train would come into conflict with other traffic on the line.

Perrow's arguments for the systemic causes of failure in high-technology systems are persuasive. They do not, however, explain some of the observations in this book. It is difficult to maintain a distinction between linear and non-linear systems. In particular, such distinctions cannot easily be used to identify situations in which accidents are 'foreseeable and avoidable' and those in which they cannot be predicted or mitigated. Even relatively simple 'systems' can fail in complex and unexpected ways. Chapter 8 reviewed the catastrophic effects of excavation activities that failed to account for soil characteristics and the knock-on effects of digging holes in particular locations. This analysis is confirmed by Petroski work on the role of failure in civil engineering [679]. Perrow counters these objections by arguing that linear processes often form part of more complex non-linear systems and hence may lead to unexpected forms of failure. This analysis does not, however, explain the lack of surprise that engineers and managers often express over the failures that affect even the most complex and tightly-coupled technological systems. This is true, for instance, of many of the adverse medical events reviewed in Chapter 14. Most iatrogenic incidents stem from recognised problems that managers lack the resources or the institutional support to address. Pressures to attain other organisational or commercial objectives stifle the concerns that are expressed in the aftermath of previous incidents. This analysis shares much in common with Sagan's view that factional interests can oppose safety measures that threaten their position within high-reliability organisations [718].

The lack of surprise that accompanys the publication of many incident reports can be explained in several ways. Theories about the causes of an incident are often formed in the aftermath of an adverse event while primary and secondary investigations are being conducted. This can lead to speculation that informs and is informed by the broadcast media. This line of analysis is unsatisfactory. The same speculation occurs both after accidents and incidents. Hence it cannot explain the differences between our observations and the surprising, catastrophic failures studied by Perrow. These differences can, however, be explained by the relatively high frequency of the adverse events that we have studied. Managers and operators are less surprised by the causes of low consequence incidents because they occur more often than the 'catastrophic' failures studied by Perrow. Further caveats can, however, be raised against this argument. In the medical domain it is possible to identify a range of high-consequence incidents whose causes are very unsurprising, including junior doctors' lack of experience and the high workloads imposed upon some staff [453, 633]. Even if a clinician never witnesses an accident, they will still be familiar with many of the causes of these events.

There is, therefore, an irony in Perrow's use of the term 'Normal Accidents'. The types of catastrophic, unexpected failures that he studies are both atypical and very rare. Coupling and complexity are, of course, major concerns. However, most accidents in high-technology systems occur through combinations of well-known problems. These unexpected combinations lend the element of surprise. In Mackay's terms, we are often familiar with the singular causes of adverse events but cannot anticipate the many causal complexes that lead to incidents or accidents [508]. I would not, however, go as far as Leveson who argues that it 'is often a matter of luck' whether the causal conditions for an incident exist [486].I would argue that the formation of causal complexes is a matter of time rather than luck. The longer organisations continue to neglect well-known singular, general causes of failure then the more likely it is that they will contribute to adverse events or near-miss incidents. Whether they develop into catastrophic accidents depends on the nature of the causal complex and the barriers that can be used to mitigate the consequences of an initial failure.

16.4 Summary

There has been a rapid growth in the number and scale of incident reporting systems. In the UK, the Ladbroke Grove rail accident stimulated a range of initiatives including the expansion of the CIRAS voluntary reporting system [194]. The Bristol Infirmary enquiry had a similar impact on UK healthcare [436]. In the United States, the Institute of Medicine report 'To Err is Human' prompted many states, as diverse as Arizona, New York and Washington, to draft bills that establish additional reporting voluntary and mandatory systems for healthcare professionals [453]. In Japan, the Maritime Labour Research Institute in Tokyo is one of several organisations that have begun to investigate alternative forms and procedures for incident analysis [553]. This builds on work in the nuclear and aviation industries. The Japanese Ministry of Public Welfare has also instructed hospitals to develop reporting systems to help reduce iatrogenic incidents. Strong claims have been made about the potential benefits of these systems. Incident reporting applications are argued to offer valuable insights into the near-miss incidents that have the potential to threaten future safety. They can also be used to elicit information about 'lessons learned' and act as an exchange for best practice [844].

This chapter has, however, argued that significant barriers must be addressed before incident reporting systems can be successfully applied within many industries. These can be summarised as follows:

1. *reporting biases.*
   There are few guarantees that all staff will contribute to a reporting system. Variations in participation rates have been observed both within working groups at the same location, as in hospital systems, and between geographical regions, for example across rail networks. Automated systems are increasingly being introduced to trigger investigations into near-miss incidents. However, some tasks cannot easily be instrumented. Many of the more specialised monitoring systems are unreliable and often provide 'false positives' that consume finite analytical resources. In consequence, it seems likely that reporting rates of less than 20-30% will be typical of many healthcare applications. These problems do not affect some reporting systems. SPAD reports provide a relatively accurate impression of the frequency of these events. However, the monitoring systems that help to detect these incidents tell us very little about events that *almost* resulted in an incident but that were narrowly averted by operator intervention.

2. *blame.*
   Some local systems enjoy good levels of participation while trusted individuals administer the scheme. Staff learn to trust the integrity of those individuals. Participation rates often fall dramatically when they are replaced [119]. This effect is clearly linked to potential contributors' concerns that they will be viewed as 'whistle blowers' either by their colleagues or by those who administer the system.

3. *analytical bias.*
   There are numerous forms of bias that can affect the analysis of incidents once they have been reported. These include author bias, judgement and hindsight bias, confirmation and frequency bias, recognition bias, political, sponsor and professional bias. This is not an exhaustive list but it illustrates the difficulty of ensuring that any investigation is not hindered by 'undue' influences. These issues are particularly important in incident reporting when many stages of an initial investigation and analysis will be performed not by an external authority but by the organisation that was directly involved in the occurrence.

4. *poor investigatory and analytical procedures.*
   Once an adverse occurrence or near miss has been reported, it can be difficult to determine what factors should be included within an investigation. This is important for theoretical reasons because it can be difficult to identify salient factors within what Mackie terms the 'causal field' [508]. Hausman also points to the problems created by 'causal asymmetry' [313]. If we know the cause then we can determine the effects. However, if we observe the effects then it can be difficult to reach firm conclusions about the multiple possible causes of those observations. These theoretical problems are exacerbated by the resource constraints that affect incident reporting. Many organisations lack both the funding and the expertise to investigate more than a single causal hypothesis. This clearly limits the value of any insights that might be obtained from the analysis of near miss incidents.

5. *inadequate risk assessments.*
   The design of safety-critical applications is typically guided by some form of risk assessment. Risk can be thought of as the product of the consequence and the likelihood of a particular failure. Incident reporting systems have been proposed as powerful means of informing risk assessments. They can provide quantitative data about the relative frequency of previous failures [453]. As we have seen, however, analytical and reporting biases undermine such statements. Similarly, the nature of 'near miss' incidents makes it very difficult to identify the 'worst plausible outcome' that might inform any decision about the consequences of a future recurrence.

6. *causation and the problems of counterfactual reasoning.*
   Many organisations have responded to the problems of analytical bias by recommending techniques that draw upon counterfactual reasoning. This style of argument takes the form; 'X is a causal factor if the incident would not have occurred if X also had not occurred'. Counterfactual reasoning is both difficult and unreliable. For example, there is often an implicit and unwarranted assumption that X occurred [313]. Similarly, Mackie's work on causal complexes suggests that investigators are likely to find many different X's for any single adverse event. The problems of counterfactual reasoning are compounded by Hausman's observations about causal asymmetries. Not only are there likely to be many X's within a causal complex, there are also likely to be many alternative causal complexes that might explain the observed effects of an incident . Unfortunately, all existing analysis techniques rely upon the subjective judgement of individual investigators to determine which of these X's are 'plausible' causes. Even the more formal, mathematically based techniques rely upon weightings or partial orders that ultimately reflect subjective assessments.

7. *classification problems.*
   Many organisations have responded to the problems of counterfactual reasoning by adopting causal taxonomies. These initiatives form part of a wider attempt to classify incidents according to a restricted range of criteria. This offers numerous benefits. In particular, the elements of the classification be used as indexing terms in relational databases. Unfortunately, field studies have shown that few safety managers know how to use these tools to accurately extract information about previous incidents. Problems also arise when the items in a database have to be manually reclassified to reflect changes in a causal taxonomy. This can be particularly onerous for national systems that hold many hundreds of thousands of records. Several prototype systems have been developed to address these problems. For instance, we are using information retrieval techniques that were originally developed for mass-market web-based applications. These approaches are the subject of on-going research and currently suffer from poor precision and recall.

8. *unrealistic expectations.*
   Many people who initiate reporting systems expect reductions in the frequency and consequence of adverse events that are unreasonable given previous experience in running these schemes. These expectations are particularly problematic given that many types of incident will not be reported to confidential systems. There can be strong organisation and cultural barriers that prevent employees from disclosing information about their friends and colleagues;

9. *reliance on reminders and quick fixes.*
   Many reporting systems lack the financial resources that are necessary to address underlying system failures. These systems are, typically, seen as a form of cost reduction rather than as a form of income generation. This separation of reporting systems from sources of investment can result in recommendations that focus narrowly on 'quick fixes'. Studies of previous systems have seen a tendency to adopt a perfective approach in which operators are urged to try harder to avoid future incidents. Such reminder statements provide dubious protection given that they must be continually reinforced if they are not to be forgotten.

10. *flawed systemic views of failure.*
    Rather than focusing on the role of individual operator error in the causes of an adverse event, attention has shifted to the 'systemic factors' that make failure more likely. For instance, Perrow has argued that the coupling and complexity of high-technology systems make it difficult to predict and prevent potential catastrophes. We would argue, however, that many operators and managers are already very familiar with the singular, general causes of adverse events. The difficulty lies in predicting how these individual factors will combine to form causal complexes.

> Reporting systems can help to address this problem. Information about previous incidents can be analysed to identify common patterns of failure. Unfortunately, the utility of this approach is compromised by the problems of analytical and reporting bias.

This is a partial list. The previous chapters in this book have mapped out a number of additional problems that complicate the development of incident reporting systems. These include the difficulty of determining whether or not a particular scheme has had any impact on safety at all. Previous chapters have also summarised the 'state of the art' in terms of the techniques that safety managers and regulators might use to address these problems. The presentation of these techniques has been driven by the use of case studies and by practical experience in applying them to previous incidents. Many of these techniques provide only partial solutions, there are no panaceas. Instead, I hope that by bringing together a wide range of material from many different industries it will be possible to learn from the successes and failures that others have experienced in the development of their reporting systems.

Bibliography

[1] A. Adams, M.A. Sasse, and P. Lunt. Making passwords secure and usable. In H. Thimbleby, B. O'Connaill, and P. Thomas, editors, *People and Computers XII: Proceedings of HCI'97*, pages 1–19, London, United Kingdom, 1997. Springer Verlag.

[2] J. Adams. *Risk.* UCL Press, London, UK, 1995.

[3] Agency for Health Care Policy and Research. CONQUEST Overview: A COmputerized Needs-oriented QUality Measurement Evaluation SysTem. Technical report, Agency for Healthcare Research and Quality, Rockville, MD, USA, 1997. http://www.ahcpr.gov/qual/conquest/conqovr1.htm.

[4] D.W. Aha and I. Watson. *Case-based Reasoning Research and Development.* Lecture Notes in Artificial Intelligence 2080. Springer Verlag, Berlin, Germany, 2001.

[5] D.W. Aha and R. Weber, editors. *Intelligent Lessons Learned Systems: Papers from the 2000 Workshop.* Technical Report WS-00-03. AAAI Press, Menlo Park, CA, USA, 2000.

[6] C. Ahlberg, C. Williamson, and B. Shneiderman. Dynamic queries for information exploration: An implementation and evaluation of graphical interfaces for drawing, exploring, and organizing. In *Proceedings of ACM SIGCHI'92 Conference on Human Factors in Computing Systems*, pages 619–626, New York, USA, 1992. ACM PRESS.

[7] P.B. Ainsworth. *Psychology, Law, and Eyewitness Testimony.* John Wiley and Sons, London, 1998.

[8] Air Accidents Investigations Branch. Report on the accident to Boeing 737-400 G-OBME near Kegworth, Leicestershire on 8th January 1989. Technical Report 4/90, Department of Transport, London, United Kingdom, 1990. http://www.open.gov.uk/aaib/gobme/gobmerep.htm.

[9] Air Accidents Investigations Branch. Report on the accident to Boeing 747-121, N739PA at Lockerbie, Dumfriesshire, Scotland on 21 December 1988. Technical Report No 2/90 (EW/C1094), Department of Transport, London, United Kingdom, 1990. http://www.open.gov.uk/aaib/n739pa.htm.

[10] Air Accidents Investigations Branch. Report on the accident to BAC One-Eleven, G-BJRT over Didcot, Oxfordshire on 10 June 1990. Technical Report 1/92, Department of Transport, London, United Kingdom, 1992. http://www.open.gov.uk/aaib/gbjrt/gbjrt.htm.

[11] Air Accidents Investigations Branch. The Civil Aviation (Investigation of Air Accidents and Incidents) Regulations 1996. Technical Report Statutory Instrument No. 2798, Department of Transport, London, United Kingdom, 1996. http://www.open.gov.uk/aaib/civact.htm.

[12] Air Accidents Investigations Branch. Report on the incident to a Boeing 737-400, G-OBMM near Daventry on 25 February 1995. Technical Report 3/96, Department of Transport, London, United Kingdom, 1996. http://www.open.gov.uk/aaib/gobmm.htm.

[13] Air Accidents Investigations Branch. Report on the accident to Aerospatiale AS 355F1 Twin Squirrel, G-CFLT Near Middlewich, Cheshire on 22 October 1996. Technical Report 4/97, Department of Transport, London, United Kingdom, 1997. http://www.open.gov.uk/aaib/gcflt/gcflt.htm.

[14] Air Accidents Investigations Branch. Report on an incident near Lambourne VOR - AIRPROX (C) : Boeing 747 and Gulfstream G IV on 3 July 1997. Technical Report 4/98, Department of Transport, London, United Kingdom, 1998. http://www.open.gov.uk/aaib/airp/airp.htm.

[15] Air Accidents Investigations Branch. Report on an incident near London Heathrow Airport on 27 August 1997 - AIRPROX (C) : Boeing 737-200 and Boeing 757. Technical Report 5/98, Department of Transport, London, United Kingdom, 1998. http://www.open.gov.uk/aaib/airp2/airp2.htm.

[16] Air Accidents Investigations Branch. Report by the Dominican Republic authorities into the accident to Boeing 757-200, G-WJAN at Puerto Plata Airfield, Dominican Republic on 1 January 1998. Technical Report 3/99, Department of Transport, London, United Kingdom, 1999. http://www.open.gov.uk/aaib/gwjan/gwjan.htm.

[17] Air Accidents Investigations Branch. Report on the accident to HS748 Series 2A, G-ATMI at Liverpool Airport on 16 August 1996. Technical Report 1/99 (EW/C96/8/8), Department of Transport, London, United Kingdom, 1999. http://www.open.gov.uk/aaib/gatmi/gatmi.htm.

[18] Air Accidents Investigations Branch. Report on the accident to Fokker F27-600 Friendship, G-CHNL Near Guernsey Airport, Channel Islands on 12 January 1999. Technical Report 2/2000 (EW/C99/1/2), Department of Transport, London, United Kingdom, 2000. http://www.open.gov.uk/aaib/gchnl/gchnl.htm.

[19] Air Accidents Investigations Branch. Report on the incidents to HS748-2A, G-BIUV and G-BGMO on approach to Isle of Man (Ronaldsway) Airport on 6 June 1998. Technical Report 1/00, Department of Transport, London, United Kingdom, 2000. http://www.open.gov.uk/aaib/gbiuv/gbiuv.htm.

[20] J. S. Aitken, P. Gray, T. Melham, and M. Thomas. Interactive theorem proving: An empirical study of user activity. *Journal of Symbolic Computation*, 25(2):263–284, February 1998.

[21] S. Aitken and T. F. Melham. An analysis of errors in interactive proof attempts. *Interacting with Computers*, 12(6):565–586, June 2000.

[22] American Forces Information Link. Fixes touted to combat friendly fire casualties. Technical report, US Department of Defense, Wshington DC, USA, 1999. http://www.defenselink.mil/news/Feb1999/n02021999_9902027.html.

[23] American Forces Information Link. Friendly fire that changed a war? Technical report, US Department of Defense, Wshington DC, USA, 1999. http://www.defenselink.mil/news/Feb1999/n02021999_9902028.html.

[24] American Institute of Chemical Engineers. Investigating process safety incidents. Technical Report Course 504, AIChE, New York, USA, 2000. http://www.aiche.org/education/cecrsdtl.asp?Number=504.

[25] M. Ammerman. *The Root Cause Analysis Handbook*. Quality Resources, New York, New York, 1998.

[26] F. Anderson. Defence Evaluation and Research Institutes: Outdated Concepts or Strategically Important Elements of the National Security Strategy of South Africa? Technical report, South African Defence College, Thaba Tshwane, Pretoria, South Africa, 1999. http://www.mil.za.

[27] J.D. Andrews and T.R. Moss. *Reliability and Risk Assessment.* Longman Scientific and Technical, Harlow, England, 1993.

[28] L.B. Andrews, C. Stocking, T. Krizek, L. Gottlieb, C. Krizek, and T. Vargish. An alternative strategy for studying adverse events in medical care. *Lancet*, pages 309–313, 1997.

[29] E. Anscombe. *Causality and Determinism.* Cambridge University Press, Cambridge, United Kingdom, 1971.

[30] ASLEF/D. N. Bennett. The Future of Safety in the Railway Industry, Letter to Rt Hon John Prescott, Deputy Prime Minister from M.D. Rix, General Secretary, ASLEF, 25 October 1999. Technical report, Associated Society of Locomotive Engineers and Firemen, London, UK, 1999. http://www.aslef.org.uk/dox/prescott.doc.

[31] ASMB Information Technology Support Team. Section 508 web alternate format. Technical report, U.S. Department of Health and Human Services, Washington, D.C., USA, 2002. http://www.hhs.gov/siteinfo/508web.html.

[32] S. Austin and G.I. Parkin. Formal methods: A survey. Technical report, Division Of Information Technology And Computing, The National Physical Laboratory, Teddington, United Kingdom, 1993.

[33] Australian Army, P. B. RETTER, Brigadier Chief of Staff. Board of Inquiry into the Death of Cadet K. P. Sperling of the South Burnett Regional Cadet Unit arising from an Incident at the Bjelke-Peterse n Dam (QLD) on 18 Nov 2000: Findings, Action and Implementation Plan. Technical report, Headquarters Training Command, Canberra, Australia, 2001. http://www.defence.gov.au/minister/2001/96180401.doc.

[34] Australian Incident Monitoring System - . GOC redevelopment. *Australian Incident Monitoring System Newsletter*, December 2000. http://www.apsf.net.au/Dec2000Newsletter.pdf.

[35] Australian Incident Monitoring System. New Features of AIMS+. *Australian Incident Monitoring System Newsletter*, December 2000. http://www.apsf.net.au/Dec2000Newsletter.pdf.

[36] Australian Incident Monitoring System. About the Australian Patient Safety Foundation. Technical report, Australian Incident Monitoring System, Australian Patient Safety Foundation, Adelaide, Australia, 2001. http://www.apsf.net.au/about.html.

[37] Australian Incident Monitoring System. APSF Coder Training Course Outline. Technical report, Australian Incident Monitoring System, Australian Patient Safety Foundation, Adelaide, Australia, 2001. http://www.apsf.net.au/coder%20training%20outline.htm.

[38] Australian Incident Monitoring System - C. Pirone. HIMS Update. *Australian Incident Monitoring System Newsletter*, March 2001. http://www.apsf.net.au/NewsletterMar01.pdf.

[39] Australian Incident Monitoring System - M. Gehrig. What's New in the Next Release of AIMS. *Australian Incident Monitoring System Newsletter*, June 2001. http://www.apsf.net.au/NewsletterJun01.pdf.

[40] Australian Incident Monitoring System - R. Parkes. e-learning the RAH experience. *Australian Incident Monitoring System Newsletter*, March 2001. http://www.apsf.net.au/NewsletterMar01.pdf.

[41] Australian Institute of Health and Welfare. Australian Institute of Health and Welfare Publications. Technical report, AIHW Media and Publishing Unit, Canberra, Australia, 2001. http://www.aihw.gov.au/publications/index.html.

[42] Australian Maritime Safety Authority. Investigation into Livestock Carrier MV Temburong on Wednesday, 27th January 1999. Technical report, AMSA, Canberra, Australia, May 1999. http://www.amsa.gov.au/amsa/pub/pub.htm.

[43] Australian Maritime Safety Authority. Organisation structure. Technical report, AMSA, Canberra, Australia, 2000. http://www.amsa.gov.au/sd/orgst.htm.

[44] Australian National Occupational Health and Safety Commission. Compendium of workers compensation statistics, Australia, 1997-98. Technical report, NOHSC, Camberra, Australia, December 1999. http://www.nohsc.gov.au/work/statistics/compstat_exec.html.

[45] Australian Therapeutic Goods Administration. Australian therapeutic device bulletin (index). Technical report, ATG, Woden ACT, Australia, 2000. http://www.health.gov.au/tga/docs/html/atdb/tdbindex.htm.

[46] Australian Transport Safety Bureau. An evaluation of the BASI-INDICATE safety program. Technical report, Department of Transport and Regional Development, Bureau of Air Safety Investigation, 1998. http://www.atsb.gov.au/aviation/pdf/indeval2.pdf.

[47] Australian Transport Safety Bureau. Collision between freight train 9784 and ballast train 9795, Ararat Victoria, 26th November 1999. Technical Report R1/2000, Rail Safety Group, 1999. http://www.atsb.gov.au/pdfrpt/ararat.pdf.

[48] Australian Transport Safety Bureau. Departmental investigation into the Engine Room Fire on board the Australian Antarctic Research and Supply Vessel Aurora Australis at the Antarctic Ice Edge, 22 July 1998. Technical Report 135, Marine Incident Investigation Unit, 1999. http://www.atsb.gov.au/marine/incident/aaustralis.cfm.

[49] Australian Transport Safety Bureau. Atsb annual review 2000. Technical report, Department of Transport and Regional Services, Canberra, Australia, 2000. http://www.atsb.gov.au/atsb/indxf/anrev.cfm.

[50] Australian Transport Safety Bureau. Incident report on Cessna Skyhawk VH-IGA and Beech B300 VH-OXF, 11 km WSW Brisbane Qld, 6 July 2000. Technical Report Occurrence Brief 200002938, Bureau of Air Safety Investigation, 2000. http://www.basi.gov.au/occurs/ob200002938.htm.

[51] Australian Transport Safety Bureau. ATSB Annual Review 2001. Technical report, Department of Transport and Regional Services, 2001. http://www.atsb.gov.au/pdf/anreview2001.pdf.

[52] Australian Transport Safety Bureau. Independent Investigation into the Shift of Cargo Aboard the General Cargo Vessel, Sea Breeze, off the West Australian port of Bunbury, 21 August 1999. Technical Report 150, Marine Safety Inspector, 2001. http://www.atsb.gov.au/marine/pdf-rep/150_sunbreeze.pdf.

[53] Australian Transport Safety Bureau. Backgrounder: Air Safety Investigations in Australia. Technical report, Department of Transport and Regional Services, Canberra, Australia, 2002. http://www.atsb.gov.au/atsb/indxf/air_invest.cfm.

[54] Australian Transport Safety Bureau. Backgrounder: The Australian Transport Safety Bureau (ATSB). Technical report, Department of Transport and Regional Services, 2002. http://www.atsb.gov.au/atsb/indxf/bground.cfm.

[55] Australian Transport Safety Bureau (Booz, Allen and Hamilton). Independent review of rail safety arrangements in Australia. Technical report, Report to the Standing Committee on Transport, Sydney, Australia, 1999. ttp://www.dotrs.gov.au/atc/rail/index.htm.

[56] Aviation Safety Reporting System. Confusion in using pre-departure clearances. Technical Report 1, NASA Ames Research Centre, California, United States of America, December 1996. http://asrs.arc.nasa.gov./operational_issues/ob_96_01.htm.

[57] Aviation Safety Reporting System. Great CRM and piloting. Technical Report 239, NASA Ames Research Centre, California, United States of America, May 1999. http://asrs.arc.nasa.gov./callback_issues/cb_239.htm.

[58] Aviation Safety Reporting System. Passenger problems. Technical report, NASA Ames Research Centre, California, United States of America, November 1999. http://asrs.arc.nasa.gov./callback_issues/cb_245.htm.

[59] Aviation Safety Reporting System. The Aviation Safety Reporting System. Technical report, NASA Ames Research Centre, California, United States of America, 2000. http://asrs.arc.nasa.gov.

[60] Aviation Safety Reporting System. The Aviation Safety Reporting System: Briefing slide 36. Technical report, NASA Ames Research Centre, California, United States of America, 2000. http://asrs.arc.nasa.gov./briefing/sld036.htm.

[61] Aviation Safety Reporting System. Insect ingestion. Technical report, NASA Ames Research Centre, California, United States of America, April 2000. http://asrs.arc.nasa.gov./callback_issues/cb_250.htm.

[62] Aviation Safety Reporting System. Unhappy landings. Technical report, NASA Ames Research Centre, California, United States of America, January 2000. http://asrs.arc.nasa.gov./callback_issues/cb_247.htm.

[63] P. Bagwell. *The railwaymen: The History of the National Union of Railwaymen.* George Allen and Unwin, London, U.K., 1963.

[64] L. Bainbridge. Analysis of verbal protocols from a process control task. In E. Edwards and F.P. Lees, editors, *The Human Operator In Process Control*, pages 146–159. Taylor And Francis, London, United Kingdom, 1974.

[65] L. Bainbridge. Ironies of automation. In J. Rasmussen, K. Duncan, and J. Leplat, editors, *New Technology And Human Error*, pages 271 – 283. J. Wiley and Sons, New York, United States of America, 1987.

[66] P. Barach and S.D. Small. Reporting and preventing medical mishaps: lessons from non-medical near miss reporting systems. *British Medical Journal*, 320(7237):759–763, 2000.

[67] R. Bareiss. *Exemplar-based knowledge acquisition: A unified approach to concept representation, classification and Learning.* Academic Press, London, United Kingdom, 1989.

[68] R. Bartlett and E. Hochstein. Air safety occurrence reporting and investigation in the Maastricht UAC. Technical report, Maastricht Air Traffic Control Centre, EUROCONTROL, Maastricht, Netherlands, 2000.

[69] R. Bastide and P. Palanque. Petri net objects for the design, validation and prototyping of user-driven interfaces. In D. Diaper, D. Gilmore, G. Cockton, and B. Shackel, editors, *Human-Computer Interaction—INTERACT'90*, pages 625–631, North Holland, Netherlands, 1990. Elsevier Science.

[70] U. Beck. *Risk Society: Towards a New Modernity.* Sage Publications, London, United Kingdom, 1992.

[71] R.K. Belew. *Finding Out About: A Cognitive perspective on Search Engine Technology and the WWW.* Cambridge University Press, Cambridge, United Kingdom, 2000.

[72] L. Benner. Accident investigation: Multilinear events sequencing methods. *Journal of Safety Research*, 7(2):67–73, 1975.

[73] L. Benner. Rating accident models and investigation methodologies. *Journal of Safety Research*, 16(3):105–126, 1985.

[74] L. Benner. Quality management for accident investigations (part i). *International Society of Air Safety Investigators' Forum*, 24(3), 1991.

[75] L. Benner. Quality management for accident investigations (part ii). *International Society of Air Safety Investigators' Forum*, 25(2), 1992.

[76] L. Benner. Review of the US Department of Energy publication on Root Cause Analysis of Performance Indicators (WP-21). Technical report, Investigation Research Roundtable and Library, http://www.iprr.org/Reviews/rca.html, 1995.

[77] T.A. Bentley and R.A. Haslam. Factors affecting postal delivery office safety performance. In S.A. Robertson, editor, *Contemporary Ergonomics*, pages 389–394. Taylor and Francis, London, United Kingdom, 1997.

[78] M. Berg and E. Goorman. The contextual nature of medical information. *Journal of Medical Informatics*, 59:51–60, 1999.

[79] P. Beynon-Davies. Human error and information systems failure: The case of the London Ambulance Service Computer-Aided Dispatch System project. *Interacting with Computers*, 11(6):699–720, 1999.

[80] V. Bignell and J. Fortune. *Understanding System Failure*. Manchester University Press, Manchester, United Kingdom, 1991.

[81] W.R. Van Biljon. Extending Petri nets for specifying man-machine dialogues. *International Journal of Man-Machine Studies*, 28(4):437–455, 1988.

[82] C.E. Billings. *Aviation Automation: The Search for a Human Centred Approach*. Lawrence Erlbaum, Mahwah, NJ, United States of America, 1997.

[83] C.E. Billings. Some hopes and concerns regarding medical event-reporting systems. *Archives of Pathology and Laboratory Medicine*, 122(3):214–215, 1998.

[84] F.E. Bird. *Management Guide to Loss Control*. Institute Press, Atlanta, GA, United States of America, 1974.

[85] F.E. Bird and G.L. Germain. Practical loss control leadership. Technical report, Institute Publishing, International Loss Control Institute, Loganville, GA, USA, 1985.

[86] A. Bobbio, L. Portinale, M. Minichino, and E. Ciancamerla. Comparing fault trees and Bayesian Networks for dependability analysis. In M. Felici, K. Kanoun, and A. Pasquini, editors, *Proceedings of SAFECOMP'99: Computer Safety, Reliability and Security*, pages 310–322, Berlin, Germany, 1999. Springer Verlag. Lecture Notes in Computer Science 1698.

[87] K. Bolte, L. Jackson, V. Roberts, and S. McComb. Accident reconstruction/simulation with event recorders. In *International Symposium on Transportation Recorders*, pages 367–369. National Transportation Safety Board, Washington DC, USA, 1999. http://www.ntsb.gov/publictn/1999/rp9901.pdf.

[88] B. Bouchard. Maximizing the benefits of recorded data. In *International Symposium on Transportation Recorders*. National Transportation Safety Board, Washington DC, USA, 1999. http://www.ntsb.gov/Events/symp_rec/proceedings/May_4/transcript_bouchard.htm.

[89] J. Bourne. Report by the comptroller general: NHS (England) summarised accounts 1995-96. Technical Report HC 127, National Audit Office Publication, London, United Kingdom, 1997.

[90] T. Bourne. The management of stress. In S.A. Robertson, editor, *Proceedings of Ergonomics'97: Contemporary Ergonomics 1997*, pages 283–288. Taylor and Francis, London, United Kingdom, 1997.

[91] C.A. Bowers, E.L. Blickensderfer, and B.B. Morgan. Air traffic control specialist team coordination. In M.W. Smolensky and E.S. Stein, editors, *Human Factors in Air Traffic Control*, pages 215–236. Academic Press, London, United Kingdom, 1998.

[92] M.H. Brandt. The next generation of FOQA programs. In *International Symposium on Transportation Recorders*, pages 29–34. National Transportation Safety Board, Washington DC, USA, 1999. http://www.ntsb.gov/publictn/1999/rp9901.pdf.

[93] T.A. Brennan, L.L. Leape LL, N.M. Laird, L. Hebert, A.R. Localio, A.G. Lawthers, J.P. Newhouse, P.C. Weiler, and H.H. Hiatt. Incidence of adverse events and negligence in hospitalized patients. results of the Harvard Medical Practice Study i. *New England Journal of Medicine*, 324(6):370–376, 1991.

[94] L. C. Briand, C. Bunse, and J. W. Daly. An experimental evaluation of quality guidelines on the maintainability of object-oriented design documents. In *Empirical Studies of Programmers: Seventh Workshop*, pages 1–19, New York, USA, 1997. ACM Press.

[95] J.C. Brigham and R.K. Bothwell. The ability of prospective jurors to estimate the accuracy of eyewitness identifications. *Law and Human Behaviour*, pages 19–30, 1983.

[96] G.J. Briscoe. MORT-based risk management. Technical Report Working Paper 28, System Safety Development Centre, E.G. & G Idaho, Inc., Idaho Falls, USA, 1991. Cited in [444].

[97] British Broadcasting Corporation. Chunnel trains 'safer than ever'. Technical report, News Staff, BBC, London, United Kingdom, 1997. http://news.bbc.co.uk/hi/english/uk/newsid

[98] British Broadcasting Corporation. Coke poisoning threat. Technical report, News Staff, BBC, London, United Kingdom, 4 February 1999. http://news.bbc.co.uk/hi/english/world/europe/newsid_272000/272440.stm#top.

[99] British Broadcasting Corporation. Rail privatisation 'confused'. Technical report, News Staff, BBC, London, United Kingdom, October 1999. http://news.bbc.co.uk/hi/english/uk_politics/newsid_480000/480950.stm.

[100] British Broadcasting Corporation. Rail summit moves forward on safety. Technical report, News Staff, BBC, London, United Kingdom, 30th November 1999. http://news.bbc.co.uk/hi/english/uk/newsid_543000/543019.stm.

[101] British Broadcasting Corporation. Safety Review After Train Disaster. Technical report, News Staff, BBC, London, United Kingdom, August 1999. http://news.bbc.co.uk/hi/english/world/south_asia/newsid_410000/410804.stm.

[102] British Broadcasting Corporation. UK Rail chiefs accused of complacency. Technical report, News Staff, BBC, London, United Kingdom, 25th October 1999. http://news.bbc.co.uk/hi/english/uk/newsid_484000/484674.stm.

[103] British Broadcasting Corporation. Kinnock can't beat euro-corruption. Technical report, News Staff, BBC, London, United Kingdom, March 2000. http://news6.thdo.bbc.co.uk/hi/english/ukstm.

[104] British Broadcasting Corporation. Mystery deepens in ferry search. Technical report, News Staff, BBC, London, United Kingdom, 2 July 2000. http://news.bbc.co.uk/hi/english/world/asia-pacific/newsid_815000/815473.stm.

[105] British Broadcasting Corporation. Plan to stop dangerous doctors. Technical report, News Staff, BBC, London, United Kingdom, 13 June 2000. http://news.bbc.co.uk/hi/english/health/newsid_788000/788805.stm.

[106] British Broadcasting Corporation. US reconsiders BNFL contract. Technical report, News Staff, BBC, London, United Kingdom, 23 March 2000. http://news2.thls.bbc.co.uk/hi/english/uk/newsid_688000/688594.stm.

[107] British Broadcasting Corporation. Bomb alert plane finally takes off. Technical report, News Staff, BBC, London, United Kingdom, 11 March 2001. http://news.bbc.co.uk/hi/english/uk/newsid_1212000/1212974.stm.

[108] British Broadcasting Corporation. Complaints about doctors soar. Technical report, News Staff, BBC, London, United Kingdom, 15 February 2001. http://news.bbc.co.uk/hi/english/health/newsid_1171000/1171412.stm.

[109] British Broadcasting Corporation. Doctors back down in whistle-blower case. Technical report, News Staff, BBC, London, United Kingdom, 2001. http://news.bbc.co.uk/hi/english/health/newsid_1470000/1470590.stm.

[110] British Broadcasting Corporation. Doctors: We're not perfect. Technical report, News Staff, BBC, London, United Kingdom, 3 July 2001. http://news.bbc.co.uk/.

[111] British Broadcasting Corporation. Milburn calls for 'end to blame'. Technical report, News Staff, BBC, London, United Kingdom, 26 June 2001. http://news.bbc.co.uk/hi/english/health/newsid_1409000/1409316.stm.

[112] British Broadcasting Corporation. Rail firm defends safety record. Technical report, News Staff, BBC, London, United Kingdom, June 2001. http://news.bbc.co.uk/hi/english/uk/scotland/newsid_1413000/1413717.stm.

[113] British Broadcasting Corporation. Inquiry urged into lorry rail crash. Technical report, News Staff, BBC, London, United Kingdom, 2002. http://news.bbc.co.uk/hi/english/uk/wales/newsid_1769000/1769983.stm.

[114] M. T. Brown. Marine voyage data recorders. In *International Symposium on Transportation Recorders*, pages 47–60. National Transportation Safety Board, Washington DC, USA, 1999. http://www.ntsb.gov/publictn/1999/rp9901.pdf.

[115] E. Brunswick. *Perception and the Representative Design of Psychological Experiments.* University of California Press, Berkley, United States of America, 1956.

[116] Bureau of Transportation Statistics. What We'll Do to Improve Safety Data. Technical report, US Department of Transportation, Washington DC, United States of America, 2002. http://www.bts.gov/sdi/section1.html.

[117] A. Burnham. A study of stop signs in Alabama and Georgia. *The Highway & Rail Safety Newsletter*, December 1994. http://ntl.bts.gov/DOCS/rmo.html.

[118] C.P. Burns. *Analysing Accidents Using Structured and Formal Methods.* PhD thesis, Department of Computing Science, University of Glasgow, 2000.

[119] D. K. Busse and D. J. Wright. Classification and analysis of incidents in complex, medical environments. *Topics in Health Information Management*, 20(4):1–11, 2000. Special Edition on Human Error and Clinical Systems.

[120] D.K. Busse and B. Holland. Implementation of critical incident reporting in a neonatal intensive care unit. *Cognition, Technology and Work*, 2001. Accepted and to appear.

[121] D.K. Busse and C.W. Johnson. Human error in an intensive care unit: A cognitive analysis of critical incidents. In J. Dixon, editor, *Proceedings of the 17th International Systems Safety Conference*, pages 138–147, Unionville, Virginia, United States of America, 1999. The Systems Safety Society.

[122] D.K. Busse and C.W. Johnson. Identification and analysis of incidents in complex, medical environments. In C.W. Johnson, editor, *Proceedings of the First Workshop on Human Error and Clinical Systems*, GAAG Technical Report G99-1, pages 101–120, Department of Computing Science, University of Glasgow, 1999. http://www.dcs.gla.ac.uk/ johnson/papers/HECS_99/.

[123] R.M.J. Byrne and S.J. Handley. Reasoning strategies for suppositional deductions. *Cognition*, pages 1–49, 1997.

[124] R.M.J. Byrne and A. Tasso. Deductive reasoning with factual, possible and counterfactual conditionals. *Memory and Cognition*, pages 726–740, 1999.

[125] P.C. Cacciabue, A. Carpignano, and C. Vivalda. A dynamic reliability technique for error assessment in man-mac hine systems. *International Journal Of Man-Machine Studies*, 38:403 – 428, 1993.

[126] P.C. Cacciabue and C. Vivalda. Dynamic methodology for evaluating human error probabilities. In G. Apostolakis, editor, *Probabilistic Safety Assessment And Management*, volume 1, pages 507–513. Elsevier, London, United Kingdom, 1991.

[127] J. Campion. Interfacing the laboratory with the real world. In J. Long and A. Whitefield, editors, *Cognitive Ergonomics And H.C.I.*, pages 35–65. Cambridge University Press, Cambridge, United Kingdom, 1989.

[128] Canadian Army Lessons Learned Centre. Analysis report - OP. ASSURANCE. Technical report, Vice-Chief of the Defence Staff, Ottawa, Canada, 1997. http://www.army.dnd.ca/allc/website/english/analysis/assur.pdf.

[129] Canadian Army Lessons Learned Centre. Common Observations and Issues: Operation Palladium Rotations Zero to Four. Technical Report Analysis 9901, Vice-Chief of the Defence Staff, Ottawa, Canada, 1999. http://www.army.dnd.ca/allc/website/english/analysis/Paladium99_e.pdf.

[130] Canadian Army, Office of the Judge Advocate General. Bill C-25: Amendments to the National Defence Act - background and amendment highlights. Technical report, Canadian Department of National Defence, Ottawa, Canada, 2001. http://www.dnd.ca/jag/hl_changeamendc25_e.html.

[131] Canadian Chief of the Defence Staff. Accident investigation and reporting. Technical Report A-GG-040-001/AG- 001, Canadian Department of National Defence, Ottawa, Canada, 2001. http://www.vcds.dnd.ca/dsafeg/safety/gsp_pp/vol1/1ch04.htm.

[132] Canadian Chief of the Defence Staff. Nuclear and ionizing radiation safety. Technical Report DAOD-4002-1, 2000-05-31, Canadian Department of National Defence, Finance and Corporate Services and Director General Nuclear Safety, Ottawa, Canada, 2001. http://www.dnd.ca/admfincs/subjects/DAOD/4002/1_e.asp.

[133] Canadian Chief of the Defence Staff. Safety and protective equipment motorcycles, motor scooters, mopeds, bicycles and snowmobiles. Technical Report CFAO 17-1, Canadian Department of National Defence, Finance and Corporate Services, Ottawa, Canada, 2001. http://www.dnd.ca/admfincs/subjects/cfao/017-01_e.asp.

[134] Canadian Department of National Defence. Initial Analysis Report - Common Issues for UN Operations in the Former Yugoslavia. Technical Report 3450-2 (ALLC), Army Lessons Learned Centre, Ottawa, Canada, 1996. http://www.army.dnd.ca/allc/website/english/analysis/obsfry.pdf.

[135] Canadian Department of National Defence. First comes safety. Technical Report Safety Digest, 1-99, Vice-Chief of the Defence Staff, Ottawa, Canada, 1999. http://www.vcds.dnd.ca/dsafeg/digest/1-99/first_e.asp.

[136] Canadian Department of National Defence. Letters to the Editor. Technical Report Safety Digest, 7-99, Vice-Chief of the Defence Staff, Ottawa, Canada, 1999. http://www.vcds.dnd.ca/dsafeg/digest/7-99/feed_e.asp.

[137] Canadian Department of National Defence. Don't 'hose' around. Technical Report Safety Digest, 8-00, Vice-Chief of the Defence Staff, Ottawa, Canada, 2000. http://www.vcds.dnd.ca/dsafeg/digest/8-00/art09_e.asp.

[138] Canadian Department of National Defence. Editorial. Technical Report Safety Digest, 5-00, Vice-Chief of the Defence Staff, Ottawa, Canada, 2000. http://www.vcds.dnd.ca/dsafeg/digest/5-00/feedback_e.asp.

[139] Canadian Department of National Defence. Go Big or Stay Home. Technical Report Safety Digest, 9-00, Vice-Chief of the Defence Staff, Ottawa, Canada, 2000. http://www.vcds.dnd.ca/dsafeg/digest/9-00/art09_e.asp.

[140] Canadian Department of National Defence. Learning from a fuel explosion. Technical Report Safety Digest, 2-00, Vice-Chief of the Defence Staff, Ottawa, Canada, 2000. http://www.vcds.dnd.ca/dsafeg/digest/2-00/fuel_e.asp.

[141] Canadian Department of National Defence. Letters to the Editor: LCdr Dave Alder. Technical Report Safety Digest, 11-00, Vice-Chief of the Defence Staff, Ottawa, Canada, 2000. http://www.vcds.dnd.ca/dsafeg/digest/11-00/art08_e.asp.

[142] Canadian Department of National Defence. Military prosecutor to appeal kipling decision. Technical Report NR-00.056, Defence News Archive, Ottawa, Canada, 2000. http://www.dnd.ca/eng/archive/2000/may00/30kipling_n_e.htm.

[143] Canadian Department of National Defence. Safety at Camp Black Bear. Technical Report Safety Digest, 2-00, Vice-Chief of the Defence Staff, Ottawa, Canada, 2000. http://www.vcds.dnd.ca/dsafeg/digest/4-00/bear_e.asp.

[144] Canadian Department of National Defence. Fluke or Just Blind Luck? Technical Report Safety Digest, 1-01, Vice-Chief of the Defence Staff, Ottawa, Canada, 2001. http://www.vcds.dnd.ca/dsafeg/digest/1-01/art05_e.asp.

[145] Canadian Department of National Defence. Letters to the Editor: Major André Bard. Technical Report Safety Digest, 6-01, Vice-Chief of the Defence Staff, Ottawa, Canada, 2001. http://www.vcds.dnd.ca/dsafeg/digest/6-01/art05_e.asp.

[146] Canadian Department of National Defence. Letters to the Editor: MCpl (ret'd) Craig Norman. Technical Report Safety Digest, 2-01, Vice-Chief of the Defence Staff, Ottawa, Canada, 2001. http://www.vcds.dnd.ca/dsafeg/digest/2-01/art07_e.asp.

[147] Canadian Department of National Defence. Gas lines and vehicles don't mix. Technical Report Safety Digest, 5-02, Vice-Chief of the Defence Staff, Ottawa, Canada, 2002. http://www.vcds.dnd.ca/dsafeg/digest/5-02/art03_e.asp.

[148] Canadian Directorate of General Safety. General accident information system. Technical report, Canadian Department of National Defence, Ottawa, Canada, 2001. http://www.vcds.dnd.ca/dsafeg/gsais-up_e.asp.

[149] Canadian Forces Individual Training and Education System (CFITES). Training and Education System: Needs Assessment. Technical report, Canadian Department of National Defence, Ottawa, Canada, 1999. http://www.dnd.ca/hr/dret/CFITES/pdf/engraph/vol2en.pdf.

[150] R. Carnap. *The Logical Foundations Of Probability.* The University of Chicago Press, Illinois, United States of America, 1962.

[151] J.M. Carroll. *The Nurnberg Funnel: Designing Minimalist Instruction For Practical Computer Skill.* MIT Press, Boston, United States of America, 1992.

[152] J. Carthy and A.F.S. Smeaton. The design of a topic tracking system. In *BCS Information Retrieval Special Interest Group Collquium*, London, UK, 2000. British Computer Society.

[153] N. Cartwright. *How the Laws of Physics Lie.* Oxford University Press, Oxford, United Kingdom, 1983.

[154] CASMET project. Casualty analysis methodology for maritime operations project, final summary report. Technical report, European Community Research and Development Information Service, Luxembourg, 1998. http://www.cordis.lu/transport/src/casmet.htm.

[155] R. Catrambone and J. M. Carroll. Learning a word processing system with training wheels and guided exploration training and advice. In *Proceedings of ACM CHI'87: Conference on Human Factors in Computing Systems and Graphics Interface*, pages 169–174, Washington, DC, USA, 1987. ACM Press.

[156] Central Labour Institute. Industrial safety division of the Indian Central Labour Institute. Technical report, CLI, Mumbai, India, 2000. http://labour.nic.in:80/dgfasli/cli/safety/safety.htm.

[157] D.B. Chaffin and G. Anderson. *Occupational Biomechanics.* J. Wiley and Sons (Wiley Interscience), New York, United States of America, 1991.

[158] M. Chalmers, R. Ingram, and C. Pfranger. Adding imageability features to information displays papers: Information visualization. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 33–39, New York, USA, 1996. ACM Press.

[159] R. Chamberlin, C. Drew, M. Patten, and R. Matchette. Ramp safety. Technical Report 8, Aviation Safety Reporting System, NASA Ames Research Centre, California, United States of America, June 1996. http://asrs.arc.nasa.gov./directline_issues/dl8_ramp.htm.

[160] Chemical Safety and Hazard Investigation Board. Investigation Report into an Explosive Manufacturing Incident (4 Deaths, 6 Injuries) at the Sierra Chemical Company, Nevada, January 7th 1998. Technical Report 98 001 I NV, CSHIB, 1998. http://www.chemsafety.gov/reports/1998/sierra_chem/.

[161] Chemical Safety and Hazard Investigation Board. Guidelines for the Board of Enquiry into the Tosco Avon Fire. Technical report, CSHIB, 1999. http://www.chemsafety.gov/1999/inv/toscoguide.htm.

[162] Chemical Safety and Hazard Investigation Board. Chemical incident reports centre. Technical report, CSHIB, 2000. http://www.chemsafety.gov/circ/.

[163] Chemical Safety and Hazard Investigation Board. Chemical safety and hazard investigation board to improve worker safety. Technical report, CSHIB, 2000. http://www.chemsafety.gov/about/who_01.htm.

[164] G. Chiola. GreatSPN users' manual. Technical report, Departmento di Informatica, Universita' delgi Studi di Tu rino, Turino, Italy, 1987.

[165] P. Chretienne. Timed Petri nets: A solution to the minimum-time-reachability problem between two states of a timed-event graph. *Journal of Systems and Software*, 6(1-2):95–101, 1986.

[166] Aeronautica Civil. Controlled Flight Into Terrain, American Airlines Flight 965, Boeing 757-223, N651AA, Near Cali, Colombia, December 20, 1995. Technical report, The Republic of Colombia, Santafe de Bogota, D.C., Colombia, 1996.

[167] H.H. Clark and S.E. Brennan. Grounding in communication. In L. B. Resnick, J. M. Levine, and S. D. Teasley, editors, *Perspectives on Socially Shared Cognition*, pages 127–149. American Psychological Association, Washington, DC, United States of America, 1991.

[168] H.H. Clark and D. Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22(1):1–39, 1986.

[169] S. Clarke. Organisational factors affecting the incident reporting of train drivers. *Work and Stress*, 12(1):6–16, 1998.

[170] S. Clarke. Safety culture on the UK railway network. *Work and Stress*, 12(3):285–292, 1998.

[171] M.R. Cohen. Why error reporting systems should be voluntary. *British Medical Journal*, 320(7237):728–729, 2000.

[172] D.J. Cole and D. Cebon. A preliminary investigation of tractor/trailer interaction in articulated vehicles. Technical Report CUED/C-MECH/TR68, Transportation Research Group, Engineering Department, Cambridge University, Cambridge, United Kingdom, July 1995. http://rage.eng.cam.ac.uk/trg/publications/abstracts/force10.html.

[173] Confidential Human Factors Incident Reporting Programme. Too much respect? *Feedback*, April 1998. http://www.chirp.co.uk/.

[174] Confidential Human Factors Incident Reporting Programme. Computer aided? *Feedback*, January 1999. http://www.chirp.co.uk/.

[175] Confidential Human Factors Incident Reporting Programme. Deaf ears? *Feedback*, July 1999. http://www.chirp.co.uk/.

[176] Confidential Human Factors Incident Reporting Programme. Excuse me Captain, but... *Feedback*, January 2000. http://www.chirp.co.uk/.

[177] Confidential Human Factors Incident Reporting Programme. Repetitive defect and sign-offs... *Feedback*, October 2000. http://www.chirp.co.uk/.

[178] Confidential Human Factors Incident Reporting Programme. Fuel fumes - not a safe option. *Feedback*, April 2002. http://www.chirp.co.uk/.

[179] Confidential Human Factors Incident Reporting Programme. Winter operations. *Feedback*, January 2002. http://www.chirp.co.uk/.

[180] E. Connors, T. Lundregan, N. Miller, and T. McEwen. Convicted by juries, exonerated by science: Case studies in the use of DNA evidence to establish innocence after trial. Technical Report NCJ 161258, National Institute of Justice, U.S. Department of Justice, Washington, DC, 1996.

[181] R. I. Cook, D. D. Woods, and C. Miller. A tale of two stories: Contrasting views of patient safety. Technical report, National Health Care Safety Council of the National Patient Safety Foundation at the American Medical Association, Chicago, Il., USA, 1998. Report from a Workshop on Assembling the Scientific Basis for Progress on Patient Safety, http://www.npsf.org/exec/toc.html.

[182] R.I. Cook and D.D. Woods. Operating at the sharp end. In M.S. Bogner, editor, *Human Error in Medicine*, pages 255–310. Lawrence Erlbaum Associates, Hillsdale, NJ, United States of America, 1994.

[183] E.P.J. Corbett and R.J. Connors. *Classical Rhetoric for the Modern Student.* Oxford University Press, Oxford, United Kingdom, 1998.

[184] Coronary Carenet. Central Cardiac Audit Database (CCAD) Project. Technical report, Society of Cardiothoracic Surgeons of Great Britain and Ireland and the British Cardiac Society, London, UK, 1997. http://ccad3.biomed.gla.ac.uk/ccad/schedule.htm.

[185] H. Van Cott. Human errors: Their causes and reduction. In M.S. Bogner, editor, *Human Error in Medicine*, pages 53–65. Lawrence Erlbaum Associates, Hillsdale, NJ, United States of America, 1994.

[186] G. Coulouris, J. Dollimore, and T. Kindberg. *Distributed Systems: Concepts and Design.* Addison Wesley, Harlow, United Kingdom, 2001. Third Edition.

[187] Council of Europe. On the control of major-accident hazards involving dangerous substances. *Official Journal of the European Communities*, L10(Directive 96/82/EC), January 1997.

[188] A.P. Cox, editor. *Risk Analysis In The Process Industries: The Report Of The International Study Group On Risk Analysis.* EFCE No. 45. Institute Of Chemical Engineers, Rugby, United Kingdom, 1985.

[189] S. Cox and T. Cox. *Safety, Systems and People.* Butterworth Heinemann, Oxford, United Kingdom, 1996.

[190] S. Cox and R. Flinn. Safety culture: Philosopher's stone or man of straw? *Work and Stress*, 12:1898–201, 1998.

[191] CRD. NHS Centre for Reviews and Dissemination. Technical report, University of York, York, UK, 2001. http://www.york.ac.uk/inst/crd/welcome.htm.

[192] J. Crow, D. Javaux, and J. Rushby. Models and mechanized methods that integrate human factors into automation design. In K. Abbott, J.-J. Speyer, and G.Boy, editors, *HCI Aero 2000: International Conference on Human-Computer Interfaces in Aeronautics*, pages 163–168, Toulouse, France, 2000. Cepadues-Editions.

[193] Cullen. *Proceedings Of The Public Enquiry Into The Piper Alpha Disaster.* The Department of Energy, London, United Kingdom, 1990.

[194] Cullen. *The Ladbroke Grove Rail Inquiry Part 1 Report.* HSE/Stationary Office, London, United Kingdom, 2001. http://www.hse.gov.uk/railway/paddrail/lgri1.pdf.

[195] R.L.P. Custer. Foresic and insurance considerations in the investigation of large loss fires. *Technology, Law and Insurance*, pages 119–124, 1998.

[196] J. Davies. Confidential Incident Reporting and Analysis System. Technical report, CIRAS, Glasgow, Scotland, UK, 2002. http://www.ciras.org.uk/index.html.

[197] J.D. Davies, L.B. Wright, E. Courtney, and H.Reid. Confidential incident reporting on UK railways: The CIRAS system. *Cognition, Technology and Work*, pages 117–125, 2000.

[198] D.M. DeJoy. Supervisor attributions and responses for multi-causal workplace accidents. *Journal of Occupational Accidents*, 9:213–223, 1987.

[199] D.M. DeJoy. An attributional model of the safety management process in industry. In S. Kumar, editor, *Advances in Industrial Ergonomics and Safety IV.* Taylor and Francis, London, UK, 1992.

[200] W.A. Dembski. *The Design Inference: Eliminating Chance Through Small Probabilities.* Cambridge University Press, Cambridge, U.K., 1998.

[201] C. Denis. Strategic plan. Technical report, Federal Aviation Administration, Washington DC, United States of America, 2000. http://www.api.faa.gov/sp00/sp2000.htm, at 11/7/2000.

[202] K. Kaur Deol, A. Sutcliffe, and N. Maiden. Towards a better understanding of usability problems with virtual environments. In M. A. Sasse and C.W. Johnson, editors, *Interact '99*, pages 544–551, Amsterdam, NL, 1999. IOS Press.

[203] Department of Energy. MORT User's Manual: For use with the Management Oversight and Risk Tree. Technical Report DOE-76/45-4-ssdc-4, Technical Research and Analysis Section, Environmental Safety and Health, U.S. Department of Energy, Washington DC, USA, 1976. http://tis.eh.doe.gov/analysis/trac/SSDC_doc/10003.txt.

[204] Department of Energy. DOE Guideline Root Cause Analysis Guidance Document. Technical Report DOE-NE-STD-1004-92, Office of Nuclear Energy and Office of Nuclear Safety Policy and Standards, U.S. Department of Energy, Washington DC, USA, 1992. http://tis.eh.doe.gov/techstds/standard/nst1004/nst1004.pdf.

[205] Department of Energy. Hazard and Barrier Analysis Guidance Document. Technical Report EH-33, Office of Operating Experience Analysis and Feedback, US Department of Energy, Washington DC, USA, 1996. http://tis.eh.doe.gov/web/tools/hazbar.pdf.

[206] Department of Energy. Argonne National Lab Cited for Nuclear Safety Violations. Technical report, Chicago Operations Office, Chicago Il, USA, 1999. http://www.ch.doe.gov/news/press/121699.htm.

[207] Department of Energy. DOE Workbook on Conducting Accident Investigations. Technical Report Revision 2, Office of the Deputy Assistant Secretary for Oversight, US Department of Energy, Washington DC, USA, 1999. http://tis.eh.doe.gov/oversight/workbook/Rev2/chpt7/chapt7.htm.

[208] Department of Energy and SCIENTECH Inc. Barrier Analysis. Technical Report SCIE-DOE-01-TRAC-29-95, Office of the Deputy Assistant Secretary for Oversight, U S Department of Energy, Washington DC, USA, 1995. http://ryker.eh.doe.gov/analysis/trac/29/trac29.html.

[209] Department of Energy and SCIENTECH Inc. Event and Causal Factor Analysis. Technical Report SCIE-DOE-01-TRAC-14-95, Office of the Deputy Assistant Secretary for Oversight, US Department of Energy, Washington DC, USA, 1995. http://ryker.eh.doe.gov/analysis/trac/14/trac14.html.

[210] Department of Transport. *Investigation into the Clapham Junction Railway Accident.* Her Majesty's Stationery Office, London, United Kingdom, 1989.

[211] K.H. Digges, P.G. Bedewi, G.T. Bahouth, N.E. Bedewi, J. Augenstein, E. Perdeck, and J. Stratton. Determination and modeling of ankle injury causation, international conference on pelvic and lower extremity injuries. Technical report, FHWA/NHTSA National Crash Analysis Center, The George Washington University and William Lehman Injury Research Center, University of Miami, Washington DC, USA, 1995. http://www.ncac.gwu.edu/archives/papers/ankle/ankle.html.

[212] Direction 2006. Direction 2006: Saving Lives Along Canada's Railways. Technical report, Transport Canada, Quebec, Canada, 2001. http://www.direction2006.com/.

[213] K. Dismukes, G. Young, and R. Sumwalt. Effective management requires a careful balancing act. Technical Report 10, Aviation Safety Reporting System, NASA Ames Research Centre, California, United States of America, December 1998. http://asrs.arc.nasa.gov./directline_issues/dl10_distract.htm.

[214] E. Dobranetski and D. Case. Pro-active use of recorded data for accident prevention. In *International Symposium on Transportation Recorders*, pages 99–120. National Transportation Safety Board, Washington DC, USA, 1999.

[215] L. Dole. On-board recorders: The black boxes of the trucking industry. In *International Symposium on Transportation Recorders*, pages 121–124. National Transportation Safety Board, Washington DC, USA, 1999.

[216] J.A. Doran and G.C. van der Graaf. Tripod-Beta: Incident Investigation and Analysis. In *Proceedings of the International Conference on Health, Safety and the Environment*, New Orleans, USA, 9-12 June 1996. Society of Petroleum Engineers.

[217] J. Dowell and W. Smith. Coordination training for distributed worksystems in emergency management. In Y. Waern, editor, *Co-operative process management - Cognition and Information Technology*, pages 147–157. Taylor and Francis, London, 1998.

[218] DTLR Press Office. 300 Deaths too many - Spellar tells conference. Technical Report News Release 2001/0437, Department of Transport, Local Government and the Regions, London, UK, 2001. http://www.press.dtlr.gov.uk.

[219] K.D. Duncan. Fault diagnosis training for advanced continuous process installations. In J. Rasmussen, K. Duncan, and J. Leplat, editors, *New Technology And Human Error*, pages 209 – 221. J. Wiley and Sons, New York, United States of America, 1987.

[220] M.D. Dunlop, C.W. Johnson, and J. Reid. Exposing the layers of information retrieval evaluation. *Interacting with Computers*, 10(3):225–237, 1998.

[221] J. Dupré. Probabilistic causality emancipated. *Midwest Studies in Philosophy*, pages 169–175, 1984. Cited in [313].

[222] M. Durkin. Digital audio recorders: Life savers, educators and vindicators. In *International Symposium on Transportation Recorders*, pages 139–144. National Transportation Safety Board, Washington DC, USA, 1999.

[223] K.W. Ellison and R. Buckhout. *Psychology and Criminal Justice*. Harper and Row, New York, 1981.

[224] R. Elmasri and S.B. Navathe. *Fundamentals of Database Systems*. Addison Wesley, Reading, MA, USA, 2000.

[225] M.R. Endsley and M.W. Smolensky. Situation awareness in air traffic control. In M.W. Smolensky and E.S. Stein, editors, *Human Factors in Air Traffic Control*, pages 108–154. Academic Press, London, United Kingdom, 1998.

[226] K.A. Ericsson and H.A. Simon. *Protocol Analysis: Verbal Reports as Data*. MIT Press, Cambridge, USA, 1985.

[227] Estonian/Finnish/Swedish Accident Investigation Commission. Joint Accident Investigation Commission of Estonia, Finland and Sweden, Final report on the Capsizing 28th September 1994 in the Baltic Sea of the Ro-Ro Ferry MV Estonia. Technical report, Estonian/Finnish/Swedish Governments, Helsinki, Finland, 1997. http://www.webandwire.com/joint.html.

[228] K. Etem and M. Patten. Communications related incidents in GA dual flight training. Technical Report 10, Aviation Safety Reporting System, NASA Ames Research Centre, California, United States of America, December 1998. http://asrs.arc.nasa.gov./directline_issues/dl10_gacom.htm.

[229] European Commission's Major Accident Hazards Bureau. Major Accident Reporting System (MARS). Technical report, Joint Research Centre, Ispra, Italy, 2000. http://mahbsrv.jrc.it/Activities-WhatIsMars.html.

[230] European Space Agency. The ESA alert system. Technical report, Product Assurance and Safety Support Division, Research and Development Arm of ESA (ESTEC), Noordwijk, The Netherlands, 1996. http://esapub.esrin.esa.it/pff/pffv6n4/ciav6n4.htm.

[231] K. Farahmand. Application of simulation modeling to emergency population evacuation. In *Proceedings of the 1997 Winter simulation conference*, pages 1181–1188, New York, United States of America, 1997. ACM Press. http://www.acm.org/pubs/citations/proceedings/simulation/307634/p1181-farahmand/.

[232] Federal Railroad Administration. Child-related railroad safety initiatives, achieving success – saving lives. Technical report, Federal Railroad Administration, Washington DC, United States of America, 1997. http://www.fra.dot.gov/safety/initives/child.htm.

[233] Federal Railroad Administration. FRA Guide for preparing accidents/incidents reports. Technical Report DOT/FRA/RRS-22 Effective: January 1997, Office of Safety, Federal Railroad Administration, Washington DC, United States of America, 1997. http://safetydata.fra.dot.gov/Objects/guide97.pdf.

[234] Federal Railroad Administration. Passenger train emergency preparedness: Notice of proposed rulemaking. Technical Report FRA Docket No. PTEP-1, Notice No. 1, RIN 2130-AA96, Office of Safety, Federal Railroad Administration, Washington DC, United States of America, 1997. http://www.fra.dot.gov/counsel/regs/uscode_working_1997/62FR8330v2.htm.

[235] Federal Railroad Administration. A Study of Supplemental Safety Systems with Whistle Bans at Highway-Rail Grade Crossings: The Spokane Experience Data Analysis. Technical report, Federal Railroad Administration, Washington DC, United States of America, 1999. Prepared by: Applied System Technologies, Inc. Rockville, MD.

[236] Federal Railroad Administration. Railroad safety inspector (qualifications and requirements). Technical Report Post GS-2121-12, Human Resources Department, Federal Railroad Administration, Washington DC, United States of America, 1999. http://www.fra.dot.gov/o/hr/railroadinspector.htm.

[237] Federal Railroad Administration. U.S. Secretary of Transportation Slater Submits Railroad Safety Bill. Technical report, Office of Public Affairs, Federal Railroad Administration/Department of Transportation, Washington DC, United States of America, 1999. http://www.dot.gov/affairs/1999/fra1899.htm.

[238] Federal Railroad Administration. U.S. Transportation Secretary Slater Announces Report Showing Dramatic Improvement In Rail Safety. Technical Report FRA 25-99, Federal Railroad Administration, Washington DC, United States of America, 1999. http://www.dot.gov/affairs/1999/fra2599.htm.

[239] Federal Railroad Administration. Mission and vision statement. Technical report, Office of the Administrator, Federal Railroad Administration, Washington DC, United States of America, 2000. http://www.fra.dot.gov/o/oad/mvtext.htm, at 11/7/2000.

[240] Federal Railroad Administration. Project Plan: 1-800 Toll-Free Emergency Notification System for Shortline Rail road Highway-Rail Crossings in the Commonwealth of Pennsylvania, A Joint Partnership Between: SEDA-COG Joint Rail Authority North Shore Railroad and Affiliated Companies Clinton County Communication Center Commonwealth of Pennsylvania Federal Railroad Administration September 20, 2000. Technical report, Federal Railroad Administration, Washington DC, United States of America, 2000.

[241] Federal Railroad Administration. Safety practices to reduce the risk of serious injury or death both to railroad employees engaged in switching operations and to the general public. Technical Report Notice of Safety Advisory 2000-03, see Federal Register, (65)213:65895, Federal Railroad Administration, Washington DC, United States of America, 2000. http://www.fra.dot.gov/safety/advisories/pdf/sa2000_3.pdf.

[242] Federal Railroad Administration. Budget Estimates, Fiscal year 2003, Federal Railroad Administration. Technical report, Federal Railroad Administration, Washington DC, United States of America, 2002. Submitted for use of Subcommittees of Appropriations, http://www.fra.dot.gov/Acquisition_and_Grant_Services/FRA_2003Budget.pdf.

[243] Federal Railroad Administration. Detailed Accident and Incident Statistics. Technical report, Office of Safety Analysis, Federal Railroad Administration, Washington DC, United States of America, 2002. http://safetydata.fra.dot.gov/Officeofsafety/default.asp.

[244] Federal Railroad Administration. Highway-Rail Crossing Web Accident Prediction System (WBAPS). Technical report, Office of Safety Analysis, Federal Railroad Administration, Washington DC, United States of America, 2002. http://safetydata.fra.dot.gov/officeofsafety/Crossing/Default.asp.

[245] Federal Railroad Administration (M. K. Coplen, Volpe National Transportation Systems Center). Compliance with Railroad Operating Rules and Corporate Culture Influences - Results of a Focus Group and Structured Interviews. Technical Report DOT/FRA/ORD-99/09, DOT-VNTSC-FRA-97-7, Office of Safety Analysis, Federal Railroad Administration, Washington DC, United States of America, 1999. http://www.fra.dot.gov/volpe/pubs/reports/ruleinfl/cvrndoc.html.

[246] J.D. Fekete and C. Plaisant. Excentric labelling: Dynamic neighbourhood labeling for data visualisation. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computer Systems*, pages 512–519, New York, United States of America, 1999. ACM Press.

[247] D. Fennell. *Investigation Into The Kings Cross Underground Fire.* Department of Transport, London, United Kingdom, 1988.

[248] T. S. Ferry. *Modern Accident Investigation and Analysis.* John Wiley and Sons Inc., London, 1988.

[249] B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read, and B. Combs. A psychometric study of attitudes towards technological risks and benefits. In P. Slovic, editor, *The Perception of Risk*, pages 80–103. ESC Publications, London, UK, 2000.

[250] J.C. Flanagan. The critical incident technique. *Psychological Bulletin*, pages 327–358, 1954.

[251] G.A. Fontenelle. The effects of task characteristics on the availability heuristic for judgements of uncertainty. Technical Report 83-1, Office of Naval Research, Washington, DC, United States of America, 1983. Jointly published with Rice University.

[252] Food and Drug Administration. Medical device user facility and manufacturer reporting, certification and registration; delegations of authority; medical device reporting procedures; final rules. *Federal Register (Rules and Regulations)*, 60(237):63577–63606, December 1995. http://www.fda.gov/cdrh/fr/fr1211t.html.

[253] Food and Drug Administration. Outbreak of Y. enterocolitica with a dairy plant. Technical report, Center for Food Safety and Applied Nutrition and the New Hampshire Division of Public Health Services Bureau of Food Protection, Washington DC, USA, 1995. http://vm.cfsan.fda.gov/ ear/NHOUTB.html.

[254] Food and Drug Administration. Vibrio vulnificus Infections Associated with Eating Raw Oysters. *Morbidity and Mortality Weekly Report*, 44(29), 1995. http://vm.cfsan.fda.gov/ mow/vulaca.html.

[255] Food and Drug Administration. Clinical impact of adverse event reporting: Postmarketing surveillance. Technical report, Department of Health and Human Services, Public Health Service, US Food and Drug Administration, Rockville, Maryland, USA, 1996. http://www.fda.gov/medwatch/articles/medcont/postmkt.htm#mw.

[256] Food and Drug Administration. FDA and the Internet advertising and promotion of medical products. Technical report, Department of Health and Human Services, FDA, Washington, DC, USA, 1996. http://www.fda.gov/opacom/morechoices/transcript1096/fdainet3.html.

[257] Food and Drug Administration. Chapter 7 recall and emergency procedures. Technical report, FDA Office of Regulatory Affairs, Washington DC, USA, 1997. http://www.fda.gov/ora/compliance_ref/rpm_new2/ch7.html.

[258] Food and Drug Administration. Medical device reporting for manufacturers. Technical report, Department of Health and Human Services, Public Health Service, US Food and Drug Administration, Rockville, Maryland, USA, March 1997. http://www.fda.gov/cdrh/manual/mdrman.html.

[259] Food and Drug Administration. National Computer Network in Place to Combat Foodborne Illness Detects and Traces E. Coli Strains Up to Five Times Faster. Technical report, U.S. Department of Health and Human Services, Washington DC, USA, 1998. http://vm.cfsan.fda.gov/ lrd/hhspulse.html.

[260] Food and Drug Administration. Transcript of the Department of Health and Human Services, Public Health Service, Food and Drug Administration, Center for Devices and Radiological Health, Stakeholders Meeting. Technical report, FDA, Washington, DC, USA, 18 August 1998 1998. http://www.fda.gov/ohrms/dockets/dockets/98N0339/98n339r/tr00001.htm.

[261] Food and Drug Administration. Workshop on minimizing medical product errors: A systems approach. Technical report, U.S. Food and Drug Administration Office of External Affairs, Center for Drug Evaluation and Research, Center for Devices and Radiological Health, Center for Biologics Evaluation and Research, Rockville, Maryland, USA, 1998. http://www.fda.gov/oc/workshops/errorsum.htm.

[262] Food and Drug Administration. Final report of a study to evaluate the feasibility and effectiveness of a sentinel reporting system for adverse event reporting of medical device use in user facilities. Technical report, Office of Surveillance and Biometrics, Center for Devices and Radiological Health, Rockville, Maryland, USA, 1999. http://www.fda.gov/cdrh/postsurv/medsunappendixa.html.

[263] Food and Drug Administration. Internet Rumors About Tampons Refuted. *FDA Consumer Magazine*, 33(2), 1999. http://www.fda.gov/fdac/departs/1999/299_upd.html, Updates and Press Releases.

[264] Food and Drug Administration. CDRH facts-on-demand index. Technical report, U.S. Food and Drug Administration, Centre for Devices and Radiological Health, Washington DC, USA, 2001. http://www.fda.gov/cdrh/dsma/dsma_ndx.html.

[265] Food and Drug Administration. Courses available for loan through the office of regulatory affairs lending library. Technical report, Office of Regulatory Affairs, Division of Human Resource Development, State Training Team, Washington, DC, USA, 2001. http://www.fda.gov/ora/training/Lending_Library2000.htm.

[266] Food and Drug Administration. FDA and Canadian Food Inspection Agency Warn Against Consuming Mislabeled Poisonous Plant Called Autumn Monkshood. Technical report, FDA, Washington DC, USA, 2001. http://www.cfsan.fda.gov/ lrd/tpplants.html.

[267] Food and Drug Administration. FDA press releases and talk papers. Technical report, U.S. Food and Drug Administration, Washington DC, USA, 2001. http://www.fda.gov/opacom/hpnews.html.

[268] Food and Drug Administration. Manufacturer and user facility device experience database (MAUDE): File formats for freedom of information releasable data. Technical report, Centre for Devices and Radiological Health, US Food and Drug Administration, Washington DC, USA, 2001. http://www.fda.gov/cdrh/maude.html.

[269] Food and Drug Administration. MedWatch: The FDA Safety Information and Adverse Event reporting Programme. Technical report, Reporting Systems Monitoring Branch (HFZ-533), Center for Devices and Radiological Health, Rockville, MD, USA, 2001. http://www.fda.gov/medwatch/index.html.

[270] Food and Drug Administration. Safety Alerts, Public Health Advisories, and Notices From CDRH. Technical report, Centre for Devices and Radiological Health, US Food and Drug Administration, Washington DC, USA, 2001. http://www.fda.gov/cdrh/safety.html.

[271] Food and Drug Administration. Title 21 Code of Federal Regulations (21 CFR Part 11) Electronic Records; Electronic Signatures Final Rule Published in the Federal Register. Technical report, U.S. Food and Drug Administration, Department of Health and Human Services, Washington DC, USA, 2001. http://www.fda.gov/ora/compliance_ref/part11/frs/background/11cfr-fr.htm.

[272] Food and Drug Administration. Manufacturer and user facility device experience database (MAUDE). Technical report, Centre for Devices and Radiological Health, US Food and Drug Administration, Washington DC, USA, 2003. http://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfMAUDE/search.CFM.

[273] Food and Drug Administration: D. Dwyer. Sending the wrong signals. *User Facility Reporting Bulletins*, 2000. http://www.fda.gov/cdrh/fusenews/ufb33.html.

[274] Food and Drug Administration: M. Meadows. Tampon Safety: TSS Now Rare, but Women Still Should Take Care. *FDA Consumer Magazine*, 32(2), 2000. http://www.fda.gov/fdac/features/2000/200_tss.html.

[275] Food and Drug Administration: M. Weick-Brady. Those codes! *User Facility Reporting Bulletins*, 1996. http://www.fda.gov/cdrh/issue18.pdf.

[276] Food and Drug Administration, Task Force on Risk Management. Managing the Risks from Medical Product Use: Creating a Risk Management Framework. Technical report, U.S. Department of Health and Human Services, Food and Drug Administration, Washington DC, USA, 1999. http://www.fda.gov/oc/tfrm/riskmanagement.pdf.

[277] J. Fortune and G. Peters. *Learning from Failure: The Systems Approach.* J. Wiley and Sons, Chichester, United Kingdom, 1997.

[278] R.B. Foster. Enhancing trust in organisations that manage risk. In M.P. Cottam, D.W. Harvey, R.P. Pape, and J. Tait, editors, *Foresight and Precaution*, volume 1, pages 3–8. Balkema, The Netherlands, 2000.

[279] H. Foushee and R.L. Helmreich. Group interaction and flight crew performance. In E. L. Wiener and D.C. Nagel, editors, *Human Factors in Aviation*, pages 189–228. Academic Press, San Diego, CA, United States of America, 1988.

[280] J.B. Freeman. *Dialectics and the Macrostructure of Arguments: A Theory of Argument Structure (Studies of Argumentation in Pragmatics and Discourse Analysis).* Foris, Dordrecht, Netherlands, 1991.

[281] D.M. Gaba. Human error in dynamic medical domains. In M.S. Bogner, editor, *Human Error in Medicine*, pages 197–224. Lawrence Erlbaum Associates, Hillsdale, NJ, United States of America, 1994.

[282] I.A.R. Galer and B.L. Yap. Ergonomics in intensive care: Applying human factors to the des ign and evaluation of a patient monitoring system. *Ergonomics*, 23(8):763 – 779, 1980.

[283] P. Galison. An accident of history. In P. Galison and A. Roland, editors, *Atmospheric Flight in the Twentieth Century*, pages 3–43, London, UK, 2000. Kluwer Academic.

[284] F. Gamst. *The Hoghead: An Industrial Ethnology of the Locomotive Engineer.* Holt, Reinhart and Winston, New York, U.S.A., 1980.

[285] A. Gawthrop. A feasability study into information retrieval tools to support searching tasks UK rail incident data. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, UK, 1999. M.Sc. Thesis.

[286] M.E. Gebicke. Military training deaths: Need to ensure that safety lessons are learned and implemented. Technical Report GAO/NSIAD-94-82, United States' General Accounting Office, Washington, DC, USA, 1994. http://www.gao.gov.

[287] M.E. Gebicke. Army equipment: Management of weapon system and equipment modification program needs improvement. Technical Report GAO/NSIAD-98-14, United States' General Accounting Office, Washington, DC, USA, 1997. http://www.gao.gov.

[288] M.E. Gebicke. Army Ranger training: Safety improvements need to be institutionalized. Technical Report GAO/NSIAD-97-29, United States' General Accounting Office, Washington, DC, USA, 1997. http://www.gao.gov.

[289] T. Gerdsmeier, P.B. Ladkin, and K. Loer. Formalising failure analysis. Technical Report Research Report RVS-Occ-97-06, Faculty of Technology, University of Bielefeld, Bielefeld, Germany, 1997. http://www.rvs.uni-bielefeld.de/publications/Reports/AMAST97.html.

[290] D.I. Gertman and H. S. Blackman. *Human Reliability and Safety Analysis Data Handbook.* John Wiley and Sons, New York, USA, 1994.

[291] W.D. Glauz, K.M. Bauer, and D.J. Migletz. Expected traffic conflict rates and their use in predicting accidents. *Transportation Research Record*, pages 1–12, 1985.

[292] L. Goodwin. The admissibility of computer-generated evidence with particular attention towards vehicle reconstruction animation evidence in court. Technical report, AIMS Research, SChEME, University of Nottingham, Nottingham, UK, 1999.

[293] R. C. Graeber. Aircrew fatigue and circadian rhythmicity. In E. L. Wiener and D.C. Nagel, editors, *Human Factors in Aviation*, pages 305–344. Academic Press, San Diego, CA, United States of America, 1988.

[294] B.A. Gran, G. Dahl, S. Eisinger, E.J. Lund, J.G. Norstrom, P. Strocka, and B.J. Ystanes. Estimating dependability of programmable systems using Bayesian belief networks. In F. Koornneef and M. van der Meulen, editors, *Computer Safety, Reliability and Security: Proceedings of 19th International Conference SAFECOMP 2000*, LNCS 1943, pages 309–320. Springer Verlag, 2000.

[295] E. Grandjean. *Fitting The Man To The Task: Occupational Ergonomics.* Taylor Francis, London, United Kingdom, 1988.

[296] P. Grice. *Studies in the Way of Words.* Harvard University Press, Cambridge, MA, United States of America, 1989.

[297] Z. Guessoum. A multi-agent simulation framework. *Transactions of the Society for Computer Simulation*, 17(1):2–11, March 2000.

[298] W. Haddon. Energy damage and the ten counter- measure strategies. *Human Factors*, 15, 1973.

[299] W. Haddon. The basic strategies for reducing damage from hazards of all kinds. *Hazard Prevention*, pages 8–12, September-October 1980.

[300] J. Hall. Opening Statement at the Pubic Hearing on the February 16, 1996 Collision Between an Amtrak Train and a MARC Commuter Train in Silver Spring, Maryland June 26, 1996. Technical report, National Transportation Safety Board, Washington, DC United States of America, 1996. http://www.ntsb.gov/speeches/former/hall/jh960626.htm.

[301] J. Hall. Remarks of Jim Hall, Chairman National Transportation Safety Board before the Senior Staff of Norfolk Southern Corporation, Roanoke, Virginia, August 14, 1998. Technical report, National Transportation Safety Board, Washington, DC United States of America, 1998. http://www.ntsb.gov/speeches/former/hall/jh980814.htm.

[302] J. Hall. Testimony of Jim Hall, Chairman National Transportation Safety Board before the Subcommittee on Aviation Committee on Transportation and Infrastructure, House of Representatives, Regarding Aviation Issues as a Result of the Crash Involving EgyptAir Flight 990. Technical report, National Transportation Safety Board, Washington, DC United States of America, 2000. http://www.ntsb.gov/speeches/jhc000411.htm.

[303] M. Hammersely and P. Atkinson. *Ethnography: Principles in Practice.* Routledge, London, UK, 2001. Second Edition.

[304] R.W. Hamming. *Introduction to Applied Numerical Analysis.* McGraw Hill, London, UK, 1971.

[305] D. Hample. The Toulmin model and the syllogism. *Journal of the American Forensic Association*, 14:1–9, 1977.

[306] M. Harris, A.P. Jagodzinski, and K.R. Greene. Human error in the context of work activity systems: A case study of knowledge based computer systems in routine use. Technical Report GAAG Technical Report G99-1, Glasgow Accident Analysis Group, Department of Computing Science, University of Glasgow, 1999.

[307] R. Harris. A handbook of rhetorical devices. Technical report, SCC, Cosa Mesa, California, 1997. http://www.sccu.edu/faculty/R_Harris/rhetoric.htm.

[308] C. Hart. The global aviation information network (GAIN): Using information proactively to improve aviation safety. Technical report, Office of System Safety, U. S. Federal Aviation Administration, Washington DC, United States of America, May 2000. http://www.asy.faa.gov/gain/What_Is_GAIN/.

[309] S.G. Hart and L.E. Staveland. Development of NASA-TLX (Task, Load Index): results of theoretical and empirical research. In P.A. Hancock and N. Meshkati, editors, *Human Mental Workload.* Elsevier Science, Amsterdam, Holland, 1988.

[310] F. Hartung and M. Kutter. Multimedia watermarking techniques. *Proceedings of the Institute of Electrical and Electronics Engineering*, 87(7):1079–1107, July 1999.

[311] J. Harvey, H. Bolam, and D. Gregory. The effectiveness of training to change safety culture and attitudes. In M.P. Cottam, D.W. Harvey, R.P. Pape, and J. Tait, editors, *Foresight and Precaution*, volume 2, pages 1143–1148, The Netherlands, 2000. Balkema.

[312] M. Hastings. A clinical review of the brain, circadian rhythms, and clock genes. *British Medical Journal*, pages 1704–1707, 1998. http://www.bmj.com/cgi/content/full/317/7174/1704.

[313] D.M. Hausman. *Causal Asymmetries*. Cambridge University Press, Cambridge, U.K., 1998.

[314] Health and Safety Commission. Health and Safety Commission discussion document on duty to investigate workplace accidents. Technical Report Press Release C48:98, Health and Safety Executive, London, United Kingdom, 1998. http://www.hse.gov.uk/press/c9848.htm, see also [316].

[315] Health and Safety Commission. Enforcement policy statement. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/action/index.htm.

[316] Health and Safety Commission. Proposal for a new duty to investigate accidents, dangerous occurrences and diseases: A consultative document. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/concdocs.

[317] Health and Safety Commission. The Southall Rail Accident Inquiry Report: HSC action plan to implement recommendations. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/hsc/south-01.htm.

[318] Health and Safety Commission. The Southall rail accident inquiry report: HSC action plan progress report to end February 2001. Technical report, Health and Safety Executive, London, United Kingdom, 2001. http://www.hse.gov.uk/hsc/uffrprt4.pdf.

[319] Health and Safety Executive. Health and safety management. Technical Report HS(G) 65, HSE, London, United Kingdom, 1991.

[320] Health and Safety Executive. The reporting of injuries, diseases and dangerous occurrences regulations 1995: Guidance for employers in the healthcare sector. Technical report, HSE, London, United Kingdom, 1995. http://www.hse.gov.uk/pubns/hsis1.htm.

[321] Health and Safety Executive. Controlling grain dust on farm. Technical Report Agriculture Information Sheet 7/96 AIS3 (revised) C100, HSE, London, United Kingdom, 1996. http://www.hse.gov.uk/pubns/ais3.htm.

[322] Health and Safety Executive. The costs to Britain of workplace accidents and work related ill health in 1995/96. Technical Report ISBN 07176 17092, HSE, London, United Kingdom, 1996. http://www.hse.gov.uk/dst/accicost.htm.

[323] Health and Safety Executive. Major hazard incidents data service (mhidas). Technical report, HSE, London, United Kingdom, 1997. http://www.hse.gov.uk/infoserv/mhidas.htm.

[324] Health and Safety Executive. Workplace injury: Comparison of Great Britain with Europe and the USA. Technical Report ISBN 07176 17092, HSE, London, United Kingdom, 1997. http://www.hse.gov.uk/hsestats/eurocomp.pdf.

[325] Health and Safety Executive. HSE enforcement and other action since the Ladbroke Grove accident. Technical report, HSE, London, United Kingdom, 1999. http://www.hse.gov.uk/railway/paddrail/enforce.htm.

[326] Health and Safety Executive. Key messages from the labour force survey (1998-99) for injury risks : Gender and age, job tenure and part time working. Technical report, HSE, London, United Kingdom, 1999. http://www.hse.gov.uk/keyart.pdf.

[327] Health and Safety Executive. Review of arrangements for standard setting and application on the main railway network : Interim report. Technical report, HSE, London, United Kingdom, 1999. http://www.hse.gov.uk/railway/standset.htm.

[328] Health and Safety Executive. Risk communication: A guide to regulatory practice. Technical report, ILGRA: Inter-Departmental Liaison Group on Risk Assessment, London, United Kingdom, 1999. http://www.hse.gov.uk/dst/risk.pdf.

[329] Health and Safety Executive. HSE charging for railway activities: Charging activities. Technical report, HSE, London, United Kingdom, 2000. http://www.hse.gov.uk/charging/railarr.htm.

[330] Health and Safety Executive. Policy, risk and science: Securing and using scientific advice. Technical report, HSE, London, United Kingdom, 2000. http://www.hse.gov.uk/research/crr_pdf/2000/crr00295.pdf.

[331] Health and Safety Executive. Safety and statistics bulletin 1999-2000. Technical report, HSE, London, United Kingdom, 2000. http://www.hse.gov.uk/hsestats/ssb9900.pdf.

[332] Health and Safety Executive. Safety statistics bulletin 1999/2000. Technical report, Strategy and Analytical Support Directorate, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/hsestats/ssb9900.pdf.

[333] Health and Safety Executive. Validation of London Underground Ltd's Railway Safety Case: Version 2. Technical report, HSE, London, United Kingdom, 2000. http://www.hse.gov.uk/railway/lul1-04.htm.

[334] Health and Safety Executive. Railway safety statistics bulletin 2000-2001. Technical report, HSE, London, United Kingdom, 2001. http://www.hse.gov.uk/railway/rihome.htm.

[335] Health and Safety Executive. Revitalising health and safety: Project plan. Technical report, HSE, London, United Kingdom, 2001. http://www.hse.gov.uk/revital/rhs.htm.

[336] Health and Safety Executive. Railways directorate: Railways industry advisory committee. Technical report, HSE, London, United Kingdom, 2002. http://www.hse.gov.uk/rail/rd_riac.htm.

[337] Health and Safety Executive, Epideamiological and Medical Statistics Unit and the Economic and Statistical Analysis Unit. Achieving the revitalising health and safety targets. Technical report, HSE, London, United Kingdom, 2001. http://www.hse.gov.uk/revital/baseline.pdf.

[338] Health and Safety Executive, Epideamiological and Medical Statistics Unit and the Economic and Statistical Analysis Unit. Achieving the revitalising health and safety targets: Statistical note on progress measurement. Technical report, HSE, London, United Kingdom, 2001. http://www.hse.gov.uk/hsestats/statnote.pdf.

[339] C. Heath, J. Hindmarsh, and P. Luff. Interaction in isolation: The dislocated world of the london underground train driver. *Sociology*, 33(3):555–575, 1999.

[340] H.W. Heinrich. *Industrial Accident Prevention*. McGraw Hill, New York, United States of America, 1931.

[341] R.L. Helmreich, R. A. Butler, W. R. Taggart, and J.A. Wilhelm. Behavioral markers in accidents and incidents: Reference list. Technical Report Technical Report 95-1, NASA/University of Texas/FAA Aerospace Crew Research Project, Austin, TX, United States of America, 1995. http://www.psy.utexas.edu/psy/helmreich/acdntlst.htm.

[342] R.L. Helmreich and A.C. Merritt. *Culture at Work in Aviation and Medicine*. Ashgate, Aldershot, United Kingdom, 1998.

[343] R.L. Helmreich, A.C. Merritt, and J.A. Wilhelm. The evolution of crew resource management in commercial aviation. *The International Journal of Aviation Psychology*, 9:19–32, 1999.

[344] R.L. Helmreich and H.-G. Schaefer. Team performance in the operating room. In M.S. Bogner, editor, *Human Error in Medicine*, pages 225–253. Lawrence Erlbaum Associates, Hillsdale, NJ, United States of America, 1994.

[345] C. Hempel. *Aspects of Scientific Explanation*. Free Press, New York, USA, 1965.

[346] K. Hendrick and L. Benner. *Investigating Accidents with Sequentially Timed and Events Plotting (STEP)*. Marcel Decker, New York, USA, 1987.

[347] Her Majesty's Government. The Merchant Shipping (Accident Reporting and Investigation) Regulations 1999. Technical Report Statutory Instrument 1999 No. 2567, Her Majesty's Stationery Office, London, United Kingdom, 1999. http://www.hmso.gov.uk/si/si1999/19992567.htm.

[348] Her Majesty's Railway Inspectorate. A report of the investigation into the collision that occurred on 8 August 1996 at Watford South junction on the line from Euston to Crewe in the Railtrack Midlands Zone. Technical report, Health and Safety Executive, London, United Kingdom, 1998.

[349] Her Majesty's Railway Inspectorate. Report on the inspection carried out by HM Railway Inspectorate during 1998/99 of the management systems in the railway industry covering signals passed at danger. Technical report, Health and Safety Executive, London, United Kingdom, 1999. http://www.hse.gov.uk/railway/spad-01.htm.

[350] Her Majesty's Railway Inspectorate. Assessment criteria for railway safety cases. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/railway/criteria/index.htm.

[351] Her Majesty's Railway Inspectorate. Internal inquiry report: Events leading up to the Ladbroke Grove rail accident on 5 October 1999. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/railway/paddrail/inq-03.htm.

[352] Her Majesty's Railway Inspectorate. Signals passed at danger (SPAD)s report for March 2000. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/railway/spad/march00.htm.

[353] Her Majesty's Railway Inspectorate. Signals passed at danger (SPAD)s report for November 2000. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/railway/spad/nov00.htm.

[354] Her Majesty's Railway Inspectorate. Signals passed at danger (SPAD)s report for October 2000. Technical report, Health and Safety Executive, London, United Kingdom, 2000. http://www.hse.gov.uk/railway/spad/oct00.htm.

[355] Her Majesty's Railway Inspectorate. Red/green zone working - a report on the progress with maximisation of green zone working on Railtrack infrastructure. Technical report, Health and Safety Executive, London, United Kingdom, 2001. http://www.hse.gov.uk/railway/rgzwrep1a.htm.

[356] Her Majesty's Railway Inspectorate (D.C. Hall) and W.S. Atkins Rail Ltd (F.P. Wiltshire). Railway safety assessment of railtrack's response to improvement notice I/RJS/991007/2 covering the top 22 signals passed most often at danger. Technical report, Health and Safety Executive, London, United Kingdom, 2002. http://www.hse.gov.uk/railway/spad/top22.pdf.

[357] Her Majesty's Stationary Office. Report of the court of inquiry and report of Mr Rotherby upon the circumstances attending the fall of a portion of the Tay bridge on the 28th December 1879. Technical Report Command C2616, HMSO, London, United Kingdom, 1880.

[358] Highways Agency. Work starts to reduce accidents at Lincolnshire blackspot. Technical Report EM/212/99, UK Department of Environment, Transport and the Regions, with Lincolnshire County Council, 1999. http://www.highways.gov.uk/news/pressrel/notices/a15/a15-15_10_1999.htm.

[359] Highways Agency. Toolkit - towards a balanced transport system. Technical report, UK Department of Environment, Transport and the Regions, 2000. http://www.highways.gov.uk/info/toolkit.

[360] K. Hodges. Structured interviews for assessing children. *Journal of Child Psychology and Psychiatry*, 34(1):49–58, 1993.

[361] E. Hollnagel. The phenotype of erroneous actions. *International Journal Of Man-Machine Studies*, 39:1–32, 1993.

[362] E. Hollnagel. *Cognitive Reliability and Analysis Method.* Elsevier Science, North Holland, Netherlands, 1998.

[363] E. Hollnagel. Looking for errors of omission and commission. *Reliability Engineering and System Safety*, 68(2):135–146, 2000.

[364] Hong Kong Marine Accident Investigation Branch. Explosion of a container on board the HK Licensed Dumb Steel Lighter 'Wong On No.1' at Stonecutters Island Public Cargo Working Area on 25th May 1999. Technical report, Multi-Lateral Policy Division, Marine Department, Hong Kong, 1999. Local Marine Inquiry No. 3 of 1999, http://www.info.gov.hk/mardep/dept/mai/elmi399.htm.

[365] Hong Kong Marine Accident Investigation Branch. Regulations and reporting requirements. Technical report, Multi-Lateral Policy Division, Marine Department, Hong Kong, 2001. http://www.info.gov.hk/mardep/dept/mai/elawr.htm.

[366] Hong Kong Marine Accident Investigation Branch. Types of marine accident investigation carried out by marine department. Technical report, Multi-Lateral Policy Division, Marine Department, Hong Kong, 2001. http://www.info.gov.hk/mardep/dept/mai/einvtype.htm.

[367] V.D. Hopkin. The impact of automation. In M.W. Smolensky and E.S. Stein, editors, *Human Factors in Air Traffic Control*, pages 391–419. Academic Press, London, United Kingdom, 1998.

[368] M. Horswill and F.P. McKenna. The effect of perceived control on risk taking. *Journal of Applied Social Psychology*, 29(2):377–391, 1999.

[369] C. Howson and P. Urbach. *Scientific Reasoning: The Bayesian Approach.* Open Court, La Salle, USA, 1993.

[370] T.W. Hoyes. Risk homeostasis theory - beyond transportational research. *Safety Science*, 17:77–89, 1994.

[371] T.W. Hoyes and A.I. Glendon. Risk homeostasis: Issues for future research. *Safety Science*, 16:19–33, 1993.

[372] P. Hudson, J. Reason, W. Wagenaar, P. Bentley, M. Primrose, and J. Visser. Tripod-Delta: Pro-active approach to enhanced safety. *Journal of Petroleum Technology*, pages 58–62, 1994.

[373] D. Huff. *How to Lie with Statistics.* Penguin Book, London, UK, 1992.

[374] J. Hughes, V. King, T. Rodden, and H. Andersen. Moving out of the control room: Ethnography in system design. In *Proceedings of CSCW'94*, pages 429–438, New York, U.S.A., 1994. ACM Press.

[375] Human Factors Group. Rail and Aviation Working Group. Technical report, Royal Aeronautical Society, London, UK, 2002. http://www.raes-hfg.com/rawgsg.htm.

[376] D. Hume. *An Enquiry Concerning Human Understanding.* Oxford Philosophical Texts, Oxford, UK, 1777 (first published). 1998 (this edition). Edited by T.L. Beauchamp.

[377] G.S. Hura and J.W. Attwood. The use of Petri nets to analyse coherent fault trees. *IEEE Transactions On Reliability*, 37(5):469–473, 1988.

[378] A. E. Hutt. *Computer Security Handbook.* John Wiley and Sons, New York, USA, 1997.

[379] N.F. Iannone and M.P. Iannone. *Supervision of Police Personnel.* Prentice Hall, New York, USA, 2000.

[380] W.R. Wells II. Friendly target. Technical report, US Naval Institute, Annapolis, Md, USA, 2001. http://www.usni.org/navalhistory/NHwells.htm.

[381] Nuclear Energy Institute. An American Success Story: What State and Local officials are saying about tranporting Used Nucelar Fuel. Technical report, NEI, Washington, DC, USA, 2000. http://www.nei.org/documents/SafeShipBrochure2.pdf.

[382] International Atomic Energy Authority. The IAEA/NEA incident reporting system: Using operational experience to improve safety. Technical report, INEA, Department of Nuclear Safety, Vienna, Austria, 2000. http://WWW.iaea.or.at/worldatom/inforesource/other/iaeanea/iaeanea-irs.html.

[383] International Civil Aviation Organisation. Accident prevention recommendations complement recent global aviation safety initiatives. Technical Report PIO 12/99, ICAO Press Information Office, Montreal, Quebec, Canada, 1999. Press release http://www.icao.int/icao/en/nr/pio9912.htm.

[384] International Civil Aviation Organisation. *Convention on International Civil Aviation.* ICAO, Montreal, Quebec, Canada, 1999 (reprinted).

[385] International Loss Control Institute. International oil and petrochemical safety rating. Technical report, ILCI, Loganville, GA, USA, 1978.

[386] International Maritime Organisation. A.853(20) Amendments to the Code for the Safe Carriage of Irradiated Nuclear Fuel, Plutonium and High-Level Radioactive Wastes in Flasks on Board Ships and Adoption of Guidelines for Developing Shipboard Emergency Plans... Technical report, Assembly of the International Maritime Organisation, London, United Kingdom, November 1997. http://www.imo.org/assembly/853854.htm.

[387] International Maritime Organisation. IMO Assembly Resolution A.849 (2.0): Code for the Investigation of Marine Casualties and Incidents. Technical report, Assembly of the International Maritime Organisation, London, United Kingdom, 1997. http://http://www.maiif.net/.

[388] International Maritime Organisation. IMO Assembly Resolution A.850 (20): The Human Element Vision, Principles and Goals of the Organisation. Technical report, Assembly of the International Maritime Organisation, London, United Kingdom, 1997. http://www.imo.org.

[389] International Maritime Organisation. Resolutions adopted by the 20th Assembly of the International Maritime Organisation. Technical report, Assembly of the International Maritime Organisation, London, United Kingdom, November 1997. http://www.imo.org/assembly/8486352.htm.

[390] Irish Department of Public Enterprise and International Risk Management Services. A review of railway safety in Ireland - implementation review, a report prepared by international risk management services on behalf of the Department of Public Enterprise. Technical Report Report No: 2081.01, Issue No: 03, Department of Public Enterprise, Dublin, Ireland, 2000. http://www.ireland.com/newspaper/special/2000/rail/.

[391] Japanese Maritime Accidents Inquiry Agency. IMO Marine Casualty and Incident Report into a Collision (a) KENRYU-MARU vs (b) HOKKAI-MARU. Technical report, Ministry of Land, Infrastructure and Transport, Tokyo, Japan, 1996. http://www.motnet.go.jp/maia.

[392] Japanese Maritime Accidents Inquiry Agency. IMO Marine Casualty and Incident Report into the Grounding of BIK DON. Technical report, Ministry of Land, Infrastructure and Transport, Tokyo, Japan, 1998. http://www.motnet.go.jp/maia.

[393] Japanese Maritime Accidents Inquiry Agency. The maritime accidents inquiry agency. Technical report, Ministry of Land, Infrastructure and Transport, Tokyo, Japan, 2001. http://www.motnet.go.jp/maia/.

[394] D. Javaux. The cognitive complexity of pilot-mode interaction. In *HCI-Aero'98: Conference on Human-machine Interaction in Aeronautics*, Montreal, Canada, 1998.

[395] M. Jeffcott. *Risk Perception for Near Miss Incidents Involving Healthcare Systems*. PhD thesis, Department of Computing Science, University of Glasgow, Glasgow, Scotland, U.K., 2002.

[396] M. Jeffcott and C.W. Johnson. On the need for risk assessment techniques to consider the role of organisational factors in information system failure within a national health service. *Cognition, Technology and Work*, 4(2):120–136, 2002.

[397] C.N. Jeffres. OSHA statement on cooperative compliance programs. Technical Report OSHA Directives, CPL 2-0.119, US Department of Labour, Washington DC, United States of America, 1998. http://www.osha.gov/media/statements/1-26-98.html.

[398] F.V. Jensen. *Bayesian Networks*. UCL Press, London, United Kingdom, 1996.

[399] F.V. Jensen, B. Chamberlain, T. Nordahl, and F. Jensen. Analysis of data conflict in HUGIN. In N.-H. Bonnisone, M. Henrion, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, pages 519–528. Elsevier, North Holland, Amsterdam, 1991.

[400] A.K. Jha, G.J. Kuperman, J.M. Teich, L. Leape, B. Shea, E. Rittenberg, E. Burdick, D.L. Seger, M. Vander Vliet, and D.W. Bates. Identifying adverse drug events: development of a computer-based monitor and comparison with chart review and stimulated voluntary report. *Journal of the American Medical Informatics Association*, 5(3):305–314, 1998.

[401] C.W. Johnson. Decision theory and safety-critical interfaces. In K. Nordby, P.H. Helmersen, D. Gilmore, and S. A. Arnesen, editors, *Interact '95*, pages 127–132. Chapman and Hall, London, United Kingdom, 1995.

[402] C.W. Johnson. Using Z to support the design of interactive, safety-critical systems software engineering. *Software Engineering Journal*, 10(2):49–60, 1995.

[403] C.W. Johnson. The evaluation of user interface design notations. In F. Bodart and J. Vanderdonk, editors, *Proceedings of the Design, Specification and Verification of Interactive Systems '96*, pages 188–206. Springer Verlag, Berlin, Germany, 1996.

[404] C.W. Johnson. Impact of working environment upon human-machine dialogues: A formal logic for the integrated specification of physical and cognitive ergonomic constraints on user interface design. *Ergonomics*, 39(3):512–530, 1996.

[405] C.W. Johnson. The epistemics of accidents. *Journal of Human Computer Systems*, 47:659–688, 1997.

[406] C.W. Johnson. Representing the impact of time on human error and systems failure. *Interacting with Computers*, 11:53–86, September 1998.

[407] C.W. Johnson. Visualizing the relationship between human error and organizational failure. In J. Dixon, editor, *Proceedings of the 17th International Systems Safety Conference*, pages 101–110, Unionville, Virginia, United States of America, 1999. The Systems Safety Society.

[408] C.W. Johnson. Why human error analysis fails to support systems development. *Interacting with Computers*, 11(5):517–524, 1999.

[409] C.W. Johnson. Don't keep reminding me: The limitations of incident reporting. In K. Abbott, J.-J. Speyer, and G.Boy, editors, *HCI Aero 2000: International Conference on Human-Computer Interfaces in Aeronautics*, pages 17–22, Toulouse, France, 2000. Cepadues-Editions.

[410] C.W. Johnson. The failure of CRM. In K. Abbott, J.-J. Speyer, and G.Boy, editors, *HCI Aero 2000: International Conference on Human-Computer Interfaces in Aeronautics*, pages 134–172, Toulouse, France, 2000. Cepadues-Editions.

[411] C.W. Johnson. Forensic software engineering. In F. Koornneef and M. van der Meulen, editors, *Computer Safety, Reliability and Security: Proceedings of 19th International Conference SAFECOMP 2000*, LNCS 1943, pages 420–430. Springer Verlag, 2000.

[412] C.W. Johnson. Proving properties of accidents. *Reliability Engineering and Systems Safety*, 67(2):175–191, 2000.

[413] C.W. Johnson. Software support for incident reporting systems in safety-critical applications. In F. Koornneef and M. van der Meulen, editors, *Computer Safety, Reliability and Security: Proceedings of 19th International Conference SAFECOMP 2000*, LNCS 1943, pages 96–106. Springer Verlag, 2000.

[414] C.W. Johnson. Using incident reporting to combat human error. In S. McDonald, Y. Waern, and G. Cockton, editors, *People and Computers XIV: Proceedings of HCI 2000*, pages 311–326, London, United Kingdom, 2000. Springer Verlag.

[415] C.W. Johnson. The London Ambulance Service, Computer Aided Dispatch System: A case study in the integration of accident reports and the constructive design of safety-critical computer systems. *Reliability Engineering and Systems Safety*, 71(3):311–326, 2001. Accepted and to appear.

[416] C.W. Johnson. Questioning the foundations of utility for quality of service in interface development. In P. Palanque and F. Paterno, editors, *Interactive Systems: Design, Specification and Verification*, pages 19–34, Berlin, Germany, 2001. Lecture Notes in Computing Science 1946.

[417] C.W. Johnson. Reasons for the failure of incident reporting in the healthcare and rail industries. In F. Redmill and T. Anderson, editors, *Components of system Safety: Proceedings of the 10th Safety-Critical Systems Symposium*, pages 31–60, Berlin, Germany, 2002. Paper from an invited talk.

[418] C.W. Johnson. The interaction between safety culture and uncertainty over device behaviour: The limitations and hazards of telemedicine. In *International Systems Safety Conference 2003*, pages 273–283, Unionville, VA, USA, 2003. International Systems Safety Society.

[419] C.W. Johnson. Newspaper and online news reporting of major accidents: Coverage in the sun, the times and bbc online of concorde flight afr4590. In C.M.Holloway C.J. Hayhurst and B. Strauch, editors, *2nd Workshop on the Investigation and Reporting of Incidents and Accidents 2003*, NASA/CP-2003-212642, pages 79–98. NASA Langley Research Centre, 2003.

[420] C.W. Johnson. Using the IEC61508 lifecycle and common requirements to guide the investigation and analysis of incidents involving electrical, electronic or programmable electronic systems. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, 2003. HSE technical report on http://www.hse.gov.uk/research/rrhtm/index.htm.

[421] C.W. Johnson and R. Botting. Using Reason's model of organisational accidents in formalising accident reports. *Cognition, Technology and Work*, 1(3):107–118, 1999.

[422] C.W. Johnson and M.D. Dunlop. Subjectivity and notions of time and value in information retrieval. *Interacting with Computers*, 10:67–75, 1998. Editorial for First Special Edition on HCI and Information Retrieval.

[423] C.W. Johnson, G. Le Galo, and M. Blaize. Guidelines for the development of occurrence reporting systems in european air traffic control. Technical report, European Organisation for Air Traffic Control (EUROCONTROL), Brussels, Belgium, 2000.

[424] C.W. Johnson and J. Gjösäther. The evaluation of user interface design notations. In F. Bodart and J. Vanderdonk, editors, *Proceedings of the Design, Specification and Verification of Interactive Systems '96*, pages 188–206. Springer Verlag, Berlin, Germany, 1996.

[425] C.W. Johnson and B. Mathers. A case study in function allocation for computer aided learning in a complex organisation. In G. Cockton, editor, *People and Computers XIV: Proceedings of HCI 2000*, Berin, Germany, 2000. Springer Verlag.

[426] C.W. Johnson, J.C. McCarthy, and P.C. Wright. Using a formal language to support natural language in accident reports. *Ergonomics*, 38(6):1265–1283, 1995.

[427] C.W. Johnson and A.J. Telford. Using formal methods to analyse human error and system failure during accident investigations. *Software Engineering Journal*, 11(6):355–365, November 1996.

[428] P. Johnson, D. Diaper, and J. Long. Task, skills and knowledge: Task analysis for knowledge based descriptions. In B. Shackel, editor, *Human-Computer Interaction - INTERACT' 84*, pages 23–27. Elsevier Science Publications, North Holland, Netherlands, 1984.

[429] W.G. Johnson. MORT the management oversight and risk tree analysis. Technical Report Technical Report SAN 8212, Atomic Energy Commission, Washington DC, USA, 1973.

[430] W.G. Johnson. *MORT Safety Assurance Systems*. Marcel Dekker, New York, USA, 1980.

[431] Joint Commission on Accrediation of Healthcare Organisations. Sentinel Events. Technical report, Sentinel Event Unit/OQM, JCAHO, Oakbrook Terrace, IL, USA, 2001. http://www.jcaho.org/sentinel/sentevnt_frm.html.

[432] D.G. Jones and M.R. Endsely. Sources of situation awareness error in aviation. *Aviation and Space Environmental Medicine*, 67(6):507–512, 1996.

[433] M.K. Junge and M.J. Giacomi. Human factors in equipment development for the Space Shuttle: A study of the general purpose work station. In R.C. Sugarman, editor, *Proceedings Of The 25th Annual Meeting Of The Human Factors Society*, pages 218–222. Human Factors Society, Santa Monica, United States of America, 1981.

[434] T. Kelly and J. McDermid. A systematic approach to safety case maintenance. In M. Felici, K. Kanoun, and A. Pasquini, editors, *Proceedings of SAFECOMP'99: Computer Safety, Reliability and Security*, pages 13–26, Berlin, Germany, 1999. Springer Verlag. Lecture Notes in Computer Science 1698.

[435] I. Kennedy. Hearing summary - 18th May 1999. Technical report, Bristol Royal Infirmary Inquiry, Bristol, United Kingdom, 2000. http://www.bristol-inquiry.org.uk/transcripts/day19.htm.

[436] I. Kennedy. *Learning from Bristol: the report of the public inquiry into children's heart surgery at the Bristol Royal Infirmary 1984 -1995*. Command Paper: CM 5207. Her Majesty's Stationary Office, London, United Kingdom, 2001. http://www.bristol-inquiry.org.uk.

[437] V. De Keyser. Temporal decision making in complex environments. In D.E. Broadbent, J. Reason, and A. Baddeley, editors, *Human Factors In Hazardous Situations*, pages 121–128. Clarendon Press, Oxford, United Kingdom, 1990.

[438] V. De Keyser. Time in ergonomics research: Recent developments and perspectives. Technical report, Faculty of Psychology, Universite Of Liege, Liege, Belgium, 1994.

[439] V. Khawani. Active train coming/second train coming sign demonstration project. Technical Report Project A-05A(1), FY 1997, with Los Angeles County Metropolitan Transportation Authority, National Academy of Sciences, 2002. See http://www4.nas.edu/trb/crp.nsf/All+Projects/TCRP+A-05A(1).

[440] D.E. Kieras, S.D. Wood, and D.E. Meyer. Predictive engineering models based on the EPIC architecture for a multi-modal high-performance human-computer interaction task. *ACM Transactions on Computer-Human Interaction*, 4(3):230–275, 1997.

[441] H.-J. Kim. Biometrics, is it a viable proposition for identity authentication and access control? *Computers and Security*, 14(3):205–214, 1995.

[442] D.J. Van Kirk. *Vehicular Accident Investigation and Reconstruction.* CRC Press, New York, USA, 2000.

[443] B. Kirwan. *A Guide to Practical Human Reliability Assessment.* Taylor and Francis, London, United Kingdom, 1994.

[444] U. Kjellen. *Prevention of Accidents Through Experience Feedback.* Taylor and Francis, London, United Kingdom, 2000.

[445] B. Klampfer and G. Grote. Integrating flight data into human factors analysis: A systems approach to incident investigation. In M. Koch and J. Dixon, editors, *17th International Systems Safety Conference*, pages 175–186, Unionville, VA, United States of America, 1999. International Systems Safety Society.

[446] R.L. Klatzy, J. Geiwitz, and S.C. Fischer. Using statistics in clinical practice: A gap between training and application. In M.S. Bogner, editor, *Human Error in Medicine*, pages 123–140. Lawrence Erlbaum Associates, Hillsdale, NJ, United States of America, 1994.

[447] K. Klein, H. Bruner, H. Holtmann, H. Rehme, J. Stolze, W. Steinhoff, and H. Wegmann. Circadian rhythms of pilots' efficiency and effects of multiple time zone travel. *Aviation Medicine*, pages 125–132, 1970.

[448] D.G. Klepacki, C.R. Morin, and R.E. Schaeffer. An investigative technique for evaluating post-accident aircraft flight control trim system configurations. In P. d'Antonio and C. Ericson, editors, *Proccedings of the 16th Annual Systems Safety Conference*, pages 336–345, Unionville, VA, United States of America, 1998. Systems Safety Society.

[449] T. Kletz. *What Went Wrong? Case Histories of Process Plant Disasters.* Gulf Publishing, Houston, TX, United States of America, 1994. 3rd Edition.

[450] D.E. Knuth. *The Art of Computer Programming: Fundamental Algorithms*, volume 1. Addison Wesley, Reading, MA, USA, 1997.

[451] D.E. Knuth. *The Art of Computer Programming: Sorting and Searching*, volume 3. Addison Wesley, Reading, MA, USA, 1998.

[452] N. Kogan and M.A. Wallach. *Risk Taking: A Study in Cognition and Personality.* Rinehart and Winston, New York, United States of America, 1964.

[453] L. Kohn, J. Corrigan, and M. Donaldson. *To Err Is Human: Building a Safer Health System.* Institute of Medicine, National Academy Press, Washington DC, United States of America, 1999. Committee on Quality of Health Care in America.

[454] J.L. Kolodner. Reconstructive memory, a computer model. *Cognitive Science*, 7:281–328, 1983.

[455] J.L. Kolodner. *Case-based Learning.* Kluwer, London, UK, 1993.

[456] F. Koornneef. *Organised Learning from Small-Scale Incidents.* PhD thesis, Technische Universiteit Delft, Delft, Netherlands, 2000.

[457] F. Koornneef and A. Hale. Using MORT to generate organisational feedback. In A. Hale, B. Wilpert, and M. Freitag, editors, *After the Event: From Accident to Organisational Learning.* Pergamon Press, Kidlington, Oxford, UK, 1997.

[458] H. Kortner and A Kjellsen. Bridging the gap between reliability centred maintenance methodology and reliability theory. In M.P. Cottam, D.W. Harvey, R.P. Pape, and J. Tait, editors, *Foresight and Precaution*, volume 1, pages 245–248, The Netherlands, 2000. Balkema.

[459] G. Kotonya and I. Sommerville. *Requirements Engineering: Processes and Techniques.* John Wiley and Sons, New York, United States of America, 1998.

[460] B. Kramer. Introducing the GRASPIN specification language SEGRAS. *Journal of Systems and Software*, 15(1):17–31, 1991.

[461] M. Kramer. The application of DesktopVR techniques to support skill transfer in safety-critical tasks. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, 1999. Masters dissertation.

[462] R. L. Kuhlman. *Professional Accident Investigation.* Institute Press, Division of International Loss Control Institute, Georgia, USA, 1977.

[463] M. P. Kurka and B. Dosa. Deliberate river-crossing operations: Focus on the fundamentals. Technical report, US Army Center for Lessons Learned and 1st Engineering Battalion (Combat) and 70th Engineering Battalion, Fort Riley, KS USA, 2000. http://call.army.mil/products/trngqtr/tq4-00/dosa.htm.

[464] L.A. Lack. Preoperative anaesthetic audit. *Bailliéres Clinical Anaesthesiology*, 4:171–184, 1990.

[465] Ladbroke Grove Rail Inquiry Secretariat. Report on Seminar - Developing an Effective Safety Culture. Technical report, LGRI, London, UK, 2000. http://www.lgri.org.uk/publicsem/seminar5.htm.

[466] Ladbroke Grove Rail Inquiry Secretariat. Report on Seminar - Employee Perspectives On Rail Safety held on Wednesday 18 October 2000. Technical report, LGRI, London, UK, 2000. http://www.lgri.org.uk/publicsem/seminar4.htm.

[467] Ladbroke Grove Rail Inquiry Secretariat. Report on Seminar - The Civil Aviation Model of Regulation. Technical report, LGRI, London, UK, 2000. http://www.lgri.org.uk/publicsem/seminar2.htm.

[468] P.B. Ladkin. On classification of factors in failures and accidents. Technical Report Document RVS-Occ-99-02, Technischen Fakultät der Universität Bielefeld, Bielefeld, Germany, 1999. http://www.rvs.uni-bielefeld.de/publications/Reports/classification.html.

[469] P.B. Ladkin. Causal reasoning about accidents. In F. Koorneef and M. van der Meulen, editors, *SAFECOMP 2000*, Lecture Notes in Computing Science No. 1943, pages 344–355. Springer Verlag, Berlin, Germany, 2000.

[470] P.B. Ladkin and K. Loer. Why-because analysis: Formal reasoning about incidents. Technical Report Document RVS-Bk-98-01, Technischen Fakultät der Universität Bielefeld, Bielefeld, Germany, 1998. http://www.rvs.uni-bielefeld.de/publications/books/WBAbook/.

[471] L. Laflamme. Age related accident ratios in assembly work: A study of female assembly workers in the Swedish automobile industry. *Safety Science*, 23:27–37, 1996.

[472] S. Lainoff. Finding human error evidence in ordinary airline event data. In M. Koch and J. Dixon, editors, *17th International Systems Safety Conference*, pages 148–152. International Systems Safety Society, Unionville, VA, United States of America, 1999.

[473] L. Lamport. The Mutual Exclusion Problem: Part I - A Theory of Interprocess Communication. *Journal of the ACM*, pages 313–326, 1986. Cited in [499].

[474] J.A. Landeweerd, I. Urlings, A. De Jong, F. Nijhuis, and L. Bouter. Risk taking tendencies amongst construction workers. *Journal of Occupational Accidents*, pages 183–196, 1990.

[475] J.C. Laprie. *Dependability: Basic Concepts and Terminology.* Springer Verlag, New York, USA, 1992.

[476] J. Law. Ladbroke Grove, or how to think about failing systems. Technical report, Department of Sociology, University of Lancaster, Lancaster, U.K., 2000. http://www.comp.lancs.ac.uk/sociology/soc055jl.html.

[477] Law Commission for England and Wales. Evidence in criminal proceedings : Hearsay and related topics. Technical Report Law Commission Report 245, Government Legal Service, London, UK, 1997. http://www.lawcom.gov.uk/library/lc245/lc245.pdf.

[478] Lawrence Livermore National Laboratory. Health and safety manual. Technical report, University of California and the U.S. Department of Energy, Livermore, California, USA, 2000. http://www.llnl.gov/es_and_h/hsm/chapter_4/chap4.html.

[479] L.L. Leape. Error in medicine. *Journal of the American Medical Association (JAMA)*, 272(23):1851–7, 1997.

[480] L.L. Leape. Editorial: Reporting of medical errors - time for a reality check. *Quality in Health Care*, 9(3):144–145, 2000.

[481] L.L. Leape. Institute of medicine medical error figures are not exaggerated. *Journal of the American Medical Association (JAMA)*, 284(1):93–95, 2000.

[482] C. C. Lebow, L. P. Sarsfield, W. L. Stanley, E. Ettedgui, and G. Henning. *Safety in the Skies: Personnel and Parties in NTSB Aviation Accident Investigations.* Institute for Civil Justice, Rand Corporation, Santa Monica, United States of America, 1999.

[483] S. Lechowicz and C. Hunt. Monitoring and managing wheel condition and loading. In *International Symposium on Transportation Recorders*, pages 205–240. National Transportation Safety Board, Washington DC, USA, 1999. http://www.ntsb.gov/publictn/1999/rp9901.pdf.

[484] A.K. Lekberg. Different approaches to incident investigation: How the analyst makes a difference. In S. Smith and B. Lewis, editors, *Proceedings of the 15th International Systems Safety Conference*, pages 178–193, Unionville, VA, United States of America, 1997. Systems Safety Society.

[485] J. Leplat. Accidents and incidents production: Methods of analysis. In J. Rasmussen, K. Duncan, and J. Leplat, editors, *New Technology and Human Error.* John Wiley and Sons, London, United Kingdom, 1987.

[486] N.G. Leveson. *Safeware: System Safety and Computers.* Addison Wesley, Reading, MA, United States of America, 1995.

[487] N.G. Leveson and C.S. Turner. An investigation of the Therac-25 accidents. *IEEE Computer*, 26(7):18–41, 1993.

[488] S.-T. Levi and A.K. Agrawala. *Real-Time Systems Design.* Mc Graw Hill, New York, United States of America, 1990.

[489] J.M. Levine and R.L. Moreland. Culture and socialization in work groups. In L. B. Resnick, J. M. Levine, and S. D. Teasley, editors, *Perspectives on Socially Shared Cognition*, pages 257–279. American Psychological Association, Washington, DC, United States of America, 1991.

[490] D. Lewis. Causation. *Journal of Philosophy*, pages 556–567, 1973.

[491] D. Lewis. *Counterfactuals.* Oxford University Press, Oxford, UK, 1973.

[492] D. Lewis. *Philosophical Papers, Vol. II.* Oxford University Press, New York, USA, 1986.

[493] M. Lind. Interpretation problems in modelling complex artifacts for diagnosis. In H. Yoshikawa and E. Hollnagel, editors, *Proceedings of the Conference on Cognitive Engineering for Process Control*, Kyoto, Japan, 12-15 November 1996.

[494] J. G. Ver Linden. Verifiers in everyday argument: An expansion of the Toulmin model. Technical report, Department of Communication, Humboldt State University, Arcata, USA, 1998. In the 84th Annual Meeting of the National Communication Association, http://www.humboldt.edu/ jgv1/ME/verifiers.html.

[495] D. S. Lindsay and J.D. Read. 'Memory work' and recovered memories of childhood sexual abuse: Scientific evidence and public, professional and personal issues. *Psychology, Public Policy and Law*, pages 846–909, 1995.

[496] P. Lipton. *Inference to the Best Explanation.* Routledge, London, 1991.

[497] B. Littlewood, P. Popov, and L. Strigini. Assessment of the reliability of fault-tolerant software: a Bayesian approach. In F. Koorneef and M. van der Meulen, editors, *SAFECOMP 2000*, Lecture Notes in Computing Science No. 1943, pages 294–308, Berlin, Germany, 2000. Springer Verlag.

[498] K. Locker. *Business and administrative communication.* Mc Graw Hill, New York, USA, 1995.

[499] K. Loer. *Towards 'Why-Because Analysis' of Failures.* PhD thesis, Technischen Fakultät der Universität Bielefeld, Bielefeld, Germany, 1998. http://www.rvs.uni-bielefeld.de/publications/Diplom/loer.ps.gz.

[500] E.F. Loftus. *Eyewitness Testimony.* Harvard University Press, Cambridge, MA, 1979.

[501] M. Lough. Design and testing an instrument for analysing a significant event. Technical report, Department of Postgraduate Medical Education, University of Glasgow, Glasgow, Sctoland, 2002.

[502] L. Love and C.W. Johnson. Using diagrams to support the analysis of system 'failure' and operator 'error'. In H. Thimbleby, B. O'Conaill, and P. Thomas, editors, *People and Computers XII: Proceedings of HCI'97*, pages 245–262. Springer Verlag, London, United Kingdom, 1997.

[503] D. Lowes and T.J.M. Bench-Capon. A prolog program to model Toulmin's argumentation schema. Technical report, Department of Computing, University of Liverpool, Liverpool, United Kingdom, 1990.

[504] C.E. Lumley, S.R. Walker, G.C. Hall, N. Staunton, and P. Grob. The under-reporting of adverse drug reactions seen in general practice. *Pharmaceutical Medicine*, 1:205–212, 1986.

[505] J.L. Lyons. Report of the inquiry board into the failure of Flight 501 of the Ariane 5 rocket. Technical report, European Space Agency, Paris, France, July 1996.

[506] K. MacCrimmon and D. Wehrung. *Taking Risks: The Management of Uncertainty.* Free Press, New York, United States of America, 1986.

[507] J.L. Mackie. The direction of causation. *Philosophical Review*, pages 441–466, 1966.

[508] J.L. Mackie. Causation and conditions. In E. Sosa and M. Tooley, editors, *Causation and Conditions*, pages 33–56. Oxford University Press, Oxford, 1993.

[509] J.D. Mackinlay, G.G. Robertson, and S.K. Card. The perspective wall: Detail and context smoothly integrated information visualization. In *Proceedings of ACM CHI'91 Conference on Human Factors in Computing Systems*, pages 173–179, New York, USA, 1991. ACM Press.

[510] R.A. Mackinnon. Rejected takeoff studies. *Aero Journal*, July 2000. http://www.boeing.com/commercial/aeromagazine/aero_11/takeoff.html.

[511] W. Macpherson. *The Stephen Lawrence Inquiry: Report of an Inquiry by Sir William Macpherson of Cluny.* Cm 4262-I. Her Majesty's Stationery Office, London, United Kingdom, 1999. http://www.official-documents.co.uk/document/cm42/4262/sli-00.htm.

[512] A. Malhotra, J. C. Thomas, J. M. Carroll, and L. A. Miller. Cognitive processes in design. *International Journal of Man-Machine Studies*, 12(2):119–140, 1980.

[513] J.G. March and H.A. Simon. *Organisations.* Wiley, New York, United States of America, 1958.

[514] Marine Accident Investigation Branch. Case 8: Complacency during routine operation. *Safety Digest*, 1 - Merchant Vessels(1), 1998. http://www.dtlr.gov.uk/maib/sd/9801/p8.htm#8.

[515] Marine Accident Investigation Branch. Case 8: Cleaner stops ship! *Safety Digest*, 1 - Merchant Vessels(02), 1999. http://www.dtlr.gov.uk/maib/sd/9902/11.htm.

[516] Marine Accident Investigation Branch. Case 1: Two recent flooding cases - vessels saved by the bilge alarm. *Safety Digest*, Fishing Vessels(Part 1: Bilge Alarms and Pumping), 2000. http://www.dtlr.gov.uk/maib/sd/00fish/04a.htm.

[517] Marine Accident Investigation Branch. Case 3: Engineer superintendent's nightmare! *Safety Digest*, 1 - Merchant Vessels(0), 2001. http://www.dtlr.gov.uk/maib/sd/0001/07.htm.

[518] Marine Accident Investigation Branch. MAIB publications. Technical report, MAIB, London, United Kingdom, 2001. http://www.dtlr.gov.uk/maib/publications.htm.

[519] Marine Accident Investigators International Forum. Guidelines to Assist Investigators in the Implementation of the International Maritime Organisation's Code for the Investigation of Marine Casualties and Incidents. Technical report, MAIIF, Canberra, Australia (currently chaired by Capt. C.W. Fitor of the Australian Transportation Safety Board), 2001. http://www.maiif.net/.

[520] Marine Incident Investigation Unit. Incidents at sea: Departmental investigation into the collision between the Panamanian bulk carrier ETERNAL WIND and the fishing vessel MELINA T, 50 miles eastward of Point Cartwright on 5 April 1998. Technical Report Incident report 131, Australian National Transportation Safety Board, Canberra, Australia, 1998. ftp://cook.dotrs.gov.au/pub/miiu/rpt131.pdf.

[521] Marine Incident Investigation Unit. Departmental investigation into the grounding of the Panama flag gerneral cargo ship NEW REACH on Heath Reef, Great Barrier Reef on 17 May 1999. Technical Report Incident report 147, Australian National Transportation Safety Board, Canberra, Australia, 1999. ftp://cook.dotrs.gov.au/pub/miiu/rpt147.pdf.

[522] R.M. Martin, K.V. Kapoor, L.V. Wilton, and R.D. Mann. Underreporting of suspected adverse drug reactions to newly marketed ('black triangle') drugs in general practice: observational study. *British Medical Journal*, 317:119–120, 1998. http://www.bmj.com/cgi/content/full/317/7151/119.

[523] Maryland Mass Transit Administration. Second train coming warning sign project. Technical report, MTA, USA, 2001. http://www.bcpl.net/ vhartsoc/stcweb.htm.

[524] B. Mathers. The use of HCI research techniques for the development of CAL within Strathclyde Fire Brigade. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, 1998. Final year dissertation.

[525] J.C. McCarthy and A.F. Monk. Measuring the quality of computer mediated communication. *Behaviour and Information Technology*, 13(5):311–319, 1994.

[526] B.H. McCormick, T.A. De Fanti, and M.D. Brown. Visualization in scientific computing. *ACM Siggraph's Computer graphics*, 21(6), 1987. Special edition of the Journal.

[527] J.A. McDermid. Issues and trends in the development of software for safety cri tical applications. Technical Report YCS 138, Department of Computer Science, University of York, York, United Kingdom, 1990.

[528] A. McDonald. The desing, implementaiton and evaluation of a web crawler for international aviation incident data. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, UK, 2002. Final Year Project report.

[529] P. McElroy. The use of information retrieval and case based reasoning tools for critical incident and accident data. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, 2000. Final year dissertation.

[530] A.G. McGill. The presentation of accident reports over the WWW. Technical report, Department of Computing Science, University of Glasgow, Glasgow, Scotland, 1998. Final year dissertation.

[531] F. McKenna. *The Railway Workers 1840-1970.* Faber and Faber, London, UK, 1980.

[532] F.P. McKenna. Do safety measures really work? An examination of risk homeostasis theory. *Ergonomics*, 28:489–498, 1985.

[533] D. McKenzie. Computer-related accidental death: An empirical exploration. *Science and Public Policy*, 21(4):233–248, 1994.

[534] John McKeown. UK AEA health, safety and the environment report 1998-1999. Technical report, UK Atomic Energy Authority, London, United Kingdon, 1999. http://www.ukaea.org.uk/sindex.htm.

[535] Medical Devices Agency. Medical devices - reporting adverse incidents and disseminating safety warnings. Technical Report MDA SN 2000(01), UK Department of Health, London, United Kingdom, January 2000. http://www.medical-devices.gov.uk/sn2000(01).htm.

[536] Medical Devices Agency. Single-use Medical Devices: Implications and Consequences of Reuse. *MDA Device Bulletin*, MDA DB2000(04), 2000. http://www.medical-devices.gov.uk/db2000(04).pdf.

[537] Medical Devices Agency. Adverse incident centre. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk/mda-aic.htm.

[538] Medical Devices Agency. Advice on the Safe Use of Bed Rails. *MDA Device Bulletin*, MDA DB2001(04), 2001. http://www.medical-devices.gov.uk/db2001(04).pdf.

[539] Medical Devices Agency. MDA annual report and accounts 2000-2001. Technical Report HC 93, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk/ann-rep2000-01.pdf.

[540] Medical Devices Agency. MDA Device Bulletins Index Page. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk/de_bulls.htm.

[541] Medical Devices Agency. MDA Device Evaluation Service Catalogue. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk/dep-intro.htm.

[542] Medical Devices Agency. MDA publications. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk/publicat.htm.

[543] Medical Devices Agency. MDA study day for nurses - infusion systems. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk/mda-studyprog2001.htm.

[544] Medical Devices Agency. Medical Devices Agency Business Plan 2001-2002. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.medical-devices.gov.uk.

[545] Medical Devices Agency. The Medical Devices Agency Annual Conference 2001 Protecting Patients - Maintaining Standards. Technical report, UK Department of Health, London, United Kingdom, 2001. http://www.neilstewartassociates.com/LX174.

[546] V.J. Mellone. TCAS II: Genie out of the bottle. Technical Report 4, Aviation Safety Reporting System, NASA Ames Research Centre, California, United States of America, June 1993. http://asrs.arc.nasa.gov./directline_issues/dl4_tcas.htm.

[547] Meteorological Office. European turbulent wake incident reporting log. Technical report, Meteorological Office/RED Scientific Limited, London, UK, 2000. http://www.meto.gov.uk/sec5/etwirl/.

[548] METROLINK (P. Hidalgo). Incidents involving trucks VS. Metrolink Trains are on an alarming rise. Technical report, METROLINK, Los Angeles, California, USA, 2000. http://www.metrolinktrains.com/news/000203_truckvtrain.htm.

[549] Ministry of Transport. Report of a public enquiry into the accident at Hixon Level Crossing on January 6th, 1968. Technical report, Ministry of Transport, London, United Kingdom, 1968.

[550] J.B. Moss and C.D. Stewart. Flamelet-based smoke properties for the field modelling of fires. *Fire Safety Journal*, 30(3):229–250, 1998.

[551] S. Moss. Critical incident management: An empirically derived computational model. *Journal of Artificial Societies and Social Simulation*, 1(4), 1998. http://www.soc.surrey.ac.uk/JASSS/1/4/1.html.

[552] S. Munson. Assessment of accident investigation methods for wildland fire-fighting incidents by case study method. Technical report, Department of Forestry, University of Montana, 1999. MSc thesis, http://www.iprr.org/3PROJ/Munpaper.html.

[553] Y. Murayama, U. Yamazaki, and M. Endo. A pilot study on marine incidents. *Journal of Japan Institute of Navigation*, 1998.

[554] R. Murray-Smith and T.A. Johansen (Eds.). *Multiple Model Approaches to Modelling and Control*. Taylor and Francis, London, United Kingdom, 1997.

[555] P. Muter and P. Maurutto. Reading and skimming from computer screens and books: The paperless office revisited? simplifying hci issues by experiment. *Behaviour and Information Technology*, 10(4):257–266, 1991.

[556] D.G. Myers. Polarising effects of social interaction. In H. Brandstatter, J.H. Davis, and G. Strocker-Kreichgauer, editors, *Group Decision Making*. Academic Press, New York, United States of America, 1982.

[557] D.C. Nagel. Human error in aviation operations. In E. L. Wiener and D.C. Nagel, editors, *Human Factors in Aviation*, pages 263–303. Academic Press, San Diego, CA, United States of America, 1988.

[558] NASA. Software documentation standard. Technical Report NASA-STD-2100-91, NASA Headquarters, Washington DC, USA, 1991. http://satc.gsfc.nasa.gov/assure/docstd.html.

[559] NASA. Software assurance standard. Technical Report NASA-STD-2201-93, NASA Headquarters, Washington DC, USA, 1992. http://satc.gsfc.nasa.gov/assure/astd.html.

[560] NASA. Software formal inspections standard. Technical Report NASA-STD-2202-93, NASA Headquarters, Washington DC, USA, 1993. http://satc.gsfc.nasa.gov/Documents/fi/std/fistdtxt.txt.

[561] NASA. Addressing the $4 billion FY 1997-2000 budget challenge: NASA's zero base review. Technical report, NASA Headquarters, Washington DC, USA, 1996. http://www.hq.nasa.gov/office/pao/ftp/budget/FY97budget/zbr.txt.

[562] NASA. Structural design and test factors of safety for spaceflight hardware. Technical Report NASA-STD-5001, NASA Headquarters, Washington DC, USA, 1996.

[563] NASA. Software safety. Technical Report NASA-STD-8719.13A, NASA Headquarters, Washington DC, USA, 1997. http://www.hq.nasa.gov/office/codeq/ns871913.htm.

[564] NASA. Mars Climate Orbiter: Mishap Investigation Board, Phase I Report. Technical report, Mars Climate Orbiter, Mishap Investigation Board, NASA Headquarters, Washington DC, USA, 1999. ftp://ftp.hq.nasa.gov/pub/pao/reports/1999/MCO_report.pdf.

[565] NASA. Computers in Spaceflight: The NASA Experience. Technical report, NASA Headquarters, Washington DC, USA, 2000. http://www.hq.nasa.gov/office/pao/History/computers/Ch4-2.html.

[566] NASA. DF88 ELT search simulation. Technical report, Mobile Aeronautics Education Laboratory, NASA Glenn Research Center, Cleveland, United States of America, 2000. http://www.lerc.nasa.gov/WWW/MAEL/ag/df88.htm.

[567] NASA. End-to-end compatibility and mission simulation testing. Technical Report Preferred reliability practices No. 1437, NASA Headquarters, Washington DC, USA, 2000. http://www.hq.nasa.gov/office/codeq/relpract/1437.pdf.

[568] NASA. Management of government safety and mission assurance surveillance functions for NASA contracts. Technical Report NASA-NPG-8735.2, NASA Headquarters, Washington DC, USA, 2000. http://nodis.hq.nasa.gov/Library/Directives/NASA-WIDE/Procedures/Program_Management/N_PG_8735_2.html.

[569] NASA. Mars Program Independent Assessment Team Report. Technical report, Mars Program Independent Assessment Team, NASA Headquarters, Washington DC, USA, 2000. http://www.jpl.nasa.gov/marsreports/mpiat_report.pdf.

[570] NASA. Report on Project Management in NASA: Phase II of the Mars Climate Orbiter Mishap Report. Technical report, Mars Climate Orbiter, Mishap Investigation Board, NASA Headquarters, Washington DC, USA, 2000. ftp://ftp.hq.nasa.gov/pub/pao/reports/2000/MCO_MIB_Report.pdf.

[571] NASA. NASA procedures and guidelines for mishap reporting, investigating and record-keeping. Technical Report NASA PG 8621.1, Safety and Risk Management Division, NASA Headquarters, Washington DC, USA, 2001. http://nodis3.gsfc.nasa.gov/displayDir.cfm?Internal$_I D = N_P R_8 621_0 01 A_{page_n ame=main}$.

[572] NASA Advisory Council. NASA federal laboratory review report. Technical report, NASA Headquarters, Washington DC, USA, 1995. http://www.hq.nasa.gov/office/fed-lab/.

[573] NASA (D. Goldin). Press release: Risk management. Technical report, NASA Headquarters, Washington DC, USA, 2000. http://www.hq.nasa.gov/office/pao/ftp/Goldin/00text/risk_management.txt.

[574] NASA (D. Goldin). "When The Best Must Do Even Better" Remarks by NASA Administrator Daniel S. Goldin At the Jet Propulsion Laboratory Pasadena, CA March 29, 2000. Technical report, NASA Headquarters, Washington DC, USA, 2000. http://www.hq.nasa.gov/office/pao/ftp/Goldin/00text/jpl_remarks.txt.

[575] NASA (Douglas Isbell). Press release: Lewis spacecraft failure board report released. Technical Report 98-109, NASA Headquarters, Washington DC, USA, 1998. http://www.hq.nasa.gov/office/pao/ftp/pressrel/1998/98-109.txt.

[576] NASA (D.W. Garrett). Press release: 1992 seen as NASA's most productive year for science discoveries. Technical Report Release: 92-228, NASA Headquarters, Washington DC, USA, 1992. http://www.hq.nasa.gov/office/pao/ftp/pressrel/1992/92-228.txt.

[577] NASA (W. Livingstone). Press release: Goldin announces initiatives to improve NASA performance. Technical Report Release: 92-154, NASA Headquarters, Washington DC, USA, 1992. http://www.hq.nasa.gov/office/pao/ftp/pressrel/1992/92-154.txt.

[578] NASA/JPL. Verification of RF Hardware Design Performance Early in the Design Phase. Technical Report Design and Test Practices for Aerospace Systems Number 1435, NASA/Jet Propulsion Laboratory, California Institute of Technology, 1996. http://techinfo.jpl.nasa.gov/www/practice/1435.pdf.

[579] NASA/JPL. Report on the loss of the Mars Polar Lander and Deep Space 2 Missions. Technical Report JPL D-18709, NASA/Jet Propulsion Laboratory, California Institute of Technology, 2000. http://www.jpl.nasa.gov/marsreports/marsreports.html.

[580] National Association of Attorneys General. Tobacco settlement announcement - news release. Technical report, NAAG, Washington, DC, United States of America, November 1998. http://www.naag.org/tobac/npr.htm.

[581] National Co-ordinating Council for Medication Error Reporting and Prevention. Taxonomy of medication errors. Technical report, NCCMERP, Rockville, Md. USA, 1998. http://www.nccmerp.org/taxo0731.pdf.

[582] National Institute of Justice. Eyewitness evidence: A guide for law enforcement. Technical Report NCJ 178240, U.S. Department of Justice, Washington, USA, 1999. http://www.ncjrs.org/pdffiles1/nij/178240.pdf.

[583] National Occupational Safety and Health Committee. Learning from safety failure. Technical report, Royal Society for the Prevention of Accident, London, UK, 2001. http://www.rospa.co.uk/.

[584] National Patient Safety Foundation. NPSF supports mandatory reporting with specific provisions to improve patient safety. Technical report, NPSF, Chicago, Il., USA, 2000. http://www.npsf.org/html/pressrel/manditrpt.htm.

[585] National Transportation Safety Board. Non-Fatal Incident Involving Cessna N3279A near BELEN, NM, June 8, 1984. Technical Report NTSB Aviation Incident Number 84A171, NTSB, Washington, DC United States of America, 1984. http://www.ntsb.gov/aviation/den/84a171.htm.

[586] National Transportation Safety Board. Fatal Accident Involving a DLM Mustang M-1 N20LL near SANTA FE, TX, July 4, 1991. Technical Report NTSB Aviation Incident Number FTW91DRA05, NTSB, Washington, DC United States of America, 1991. http://www.ntsb.gov/aviation/ftw/91ra05.htm.

[587] National Transportation Safety Board. Fatal Accident Involving Piper PA39 near Viburnam, Missouri, February 18, 1994. Technical Report NTSB Aviation Incident Number CHI94DCA01, NTSB, Washington, DC United States of America, 1994. http://www.ntsb.gov/aviation/chi/lnarr%5F94ca01.htm.

[588] National Transportation Safety Board. Pipeline accident report UGI Utilities, INC., Natural Gas Distribution Pipeline Explosion and Fire Allentown, Pennsylvania, June 9, 1994. Technical Report NTSB Pipeline Accident Number: NTSB/PAR-96/01 (PB96-916501), NTSB, Washington, DC United States of America, 1994. http://www.ntsb.gov/Publictn/1996/PAR9601.pdf.

[589] National Transportation Safety Board. Presumed Fatal Accident Involving Piper PA-28-161 N5916V. Technical Report NTSB Aviation Incident Number MIA95FAMS1, NTSB, Washington, DC United States of America, 1994. http://www.ntsb.gov/aviation/mia/95AMS1.htm.

[590] National Transportation Safety Board. Air traffic control equipment outages: Special investigation report. Technical Report PB96-91700 (NTSB/SIR-96/01), NTSB, Washington, DC United States of America, 1996. http://www.ntsb.gov/publictn/1996/sir9601.pdf.

[591] National Transportation Safety Board. In-flight fire/emergency landing Federal Express flight 1406 Douglas DC-1 0-10, N68055 Newburgh, New York, September 5, 1996. Technical Report NTSB/AAR-98/03, NTSB, Washington, DC United States of America, 1996. http://www.ntsb.gov/Publictn/1998/AAR9803.pdf.

[592] National Transportation Safety Board. Natural Gas Pipeline Rupture and Fire During Dredging of Tiger Pass, Louisiana October 23, 1996. Technical Report NTSB Pipeline Accident Number: PAR-98/01/SUM (PB98-916501), NTSB, Washington, DC United States of America, 1996. http://www.ntsb.gov/publictn/1998/par9801s.htm.

[593] National Transportation Safety Board. Release of hazardous liquid from Marathon gasoline pipe near Gramercy, Louisiana May 23, 1996. Technical Report NTSB Pipeline Accident Number: DCA-96-MP-004, NTSB, Washington, DC United States of America, 1996. http://www.ntsb.gov/publictn/1998/pab9801.htm.

[594] National Transportation Safety Board. Marine accident report grounding of the Panamanian Passenger Ship Royal Majesty on Rose and Crown Shoal near Nantucket, Massachusetts, June 10, 1995. Technical Report NTSB/MAR-97/0l, NTSB, Washington, DC United States of America, 1997. http://www.ntsb.gov/Publictn/1997/MAR9701.pdf.

[595] National Transportation Safety Board. Non-Fatal Incident Involving Gowan LONG-EZ near ROANOKE, Texas on 5th March 1997. Technical Report NTSB Aviation Incident Number FTW97FA173, NTSB, Washington, DC United States of America, 1997. http://www.ntsb.gov/aviation/ftw/lnarr%5F97a173.htm.

[596] National Transportation Safety Board. Railroad accident report collision and derailment of Maryland Rail Commuter MARC train 286 and National Railroad Passenger Corporation AMTRAK train 29 near Silver Spring, Maryland on February 16, 1996. Technical Report PB97-916302 (NTSB/RAR-97/02), NTSB, Washington, DC United States of America, 1997. http://www.ntsb.gov/Publictn/1997/RAR9702.pdf.

[597] National Transportation Safety Board. Rupture and fire involving natural gas, Indianapolis, Indiana, July 21 1997. Technical Report NTSB Pipeline Accident Number: DCA-97-FP-005, NTSB, Washington, DC United States of America, 1997. http://www.ntsb.gov/publictn/1999/pab9902.htm.

[598] National Transportation Safety Board. Safety study: Protecting excavation public safety through damage prevention. Technical Report NTSB/SS-97/0l (PB97-917003), NTSB, Washington, DC United States of America, 1997. http://www.ntsb.gov/Publictn/1997/SS9701.pdf.

[599] National Transportation Safety Board. Highway accident report: multiple vehicle crossover accident, Slinger, Wisconsin, February 12, 1997. Technical Report NTSB/HAR-98/01 (PB98-916203), NTSB, Washington, DC United States of America, 1998. http://www.ntsb.gov/publictn/1998/har9801.pdf.

[600] National Transportation Safety Board. In-Flight Icing Encounter and Uncontrolled Collision with Terrain COMAIR Flight 3272, Embraer EMB-120RT, N265CA Monroe, Michigan, January 9, 1997. Technical Report NTSB Aviation Accident Report 98-04, NTSB, Washington, DC United States of America, 1998. http://www.ntsb.gov/Publictn/1998/AAR9804.pdf.

[601] National Transportation Safety Board. Marine accident report: Fire aboard the tug Scandia and the subsequent grounding of the tug and the tank barge North Cape on Moonstone Beach, South Kingston, Rhode Island January 19, 1996. Technical Report NTSB/MAR-98/03 (PB98-916403), NTSB, Washington, DC United States of America, 1998. http://www.ntsb.gov/publictn/1998/mar9803.pdf.

[602] National Transportation Safety Board. Pipeline accident report: Natural gas pipeline rupture and subsequent explosion St. Cloud, Minnesota, December 11, 1998. Technical Report NTSB Pipeline Accident Number: NTSB/PAR-00/01 (PB2000-916501), NTSB, Washington, DC United States of America, 1998. http://www.ntsb.gov/Publictn/2000/PAR0001.pdf.

[603] National Transportation Safety Board. Railroad accident report: Derailment of AMTRAK train 4, Southwest Chief, on the Burlington Northern Santa Fe Railway near Kingman, Arizona August 9, 1997. Technical Report NTSB/RAR-98/03 (PB98-916303), NTSB, Washington, DC United States of America, 1998. http://www.ntsb.gov/publictn/1998/rar9803.pdf.

[604] National Transportation Safety Board. We are all safer: NTSB inspired improvements in transportation safety. Technical report, NTSB, Washington, DC United States of America, 1998. http://www.ntsb.gov/publictn/1998/SR9801.htm.

[605] National Transportation Safety Board. Accident investigation docket, images from the board meeting on usair flight 427, march 23, 1999. Technical report, NTSB, Washington, DC, United States of America, 1999. http://www.ntsb.gov/events/usair427/images.htm, see also [609].

[606] National Transportation Safety Board. Fire on board the Liberian passenger ship Ecstasy Miami, Florida July 20, 1998. Technical Report PB2001-916401, NTSB/MAR-01/01, NTSB, Washington, DC United States of America, 1999. http://www.ntsb.gov/Publictn/2001/MAR0101.htm.

[607] National Transportation Safety Board. Highway special investigation report: Selective motorcoach issues. Technical Report NTSB/SIR-99/01, NTSB, Washington, DC United States of America, 1999. http://www.ntsb.gov/publictn/1999/SIR9901.pdf.

[608] National Transportation Safety Board. Safety Study: Evaluation of U.S. Department of Transportation Efforts in the 1990s to Address Operator Fatigue. Technical Report Safety Report NTSB/SR-99/01 May 1999 PB99-917002 Notation 7155, NTSB, Washington, DC United States of America, 1999. http://www.ntsb.gov/Publictn/1999/SR9901.pdf.

[609] National Transportation Safety Board. Uncontrolled descent and collision with terrain USAir Flight 427, Boeing 737-300, N513AU near Aliquippa, Pennsylvania September 8, 1994. Technical Report NTSB Number AAR-99/01, NTSB, Washington, DC, United States of America, 1999. http://www.ntsb.gov/Publictn/1999/AAR9901.pdf, see also [605].

[610] National Transportation Safety Board. Effectiveness of commercial driver oversight programs (video presentations) public hearing, January 20-21, New Orleans, Lousiana. Technical Report NTSB Accident No. HWY-99-M-H017, NTSB, Washington, DC United States of America, 2000. http://www.ntsb.gov/events/2000/comm_driver/animations.htm.

[611] National Transportation Safety Board. NTSB Selects George Washington University Site for its Training Academy. Technical report, NTSB, Washington, DC, United States of America, 2000. http://www.ntsb.gov/pressrel/2000/001113.htm.

[612] National Transportation Safety Board. Railroad accident report derailment of Burlington Northern and Santa Fe Railway Company intermodal freight train S-CHILAC1-31, Crisfield, Kansas September 2, 1998. Technical Report NTSB/RAR-00/01 (PB2000-916301), NTSB, Washington, DC United States of America, 2000. http://www.ntsb.gov/Publictn/2000/RAR0001.pdf.

[613] National Transportation Safety Board. Safety Recommendations to the International Council of Cruise Lines and Cruise Line Companies Regarding Fires on Board Passenger Ships, Public meeting of July 11, 2000 (subject to editing). Technical report, NTSB, Washington, DC, United States of America, 2000. http://www.ntsb.gov/pressrel/2000/000711.htm.

[614] National Transportation Safety Board. Statement of Jim Hall, Chairman National Transportation Safety Board at Closing of NTSB Board Meeting Final Report Crash of TWA Flight 800 August 23, 2000. Technical report, NTSB, Washington, DC, United States of America, 2000. http://www.ntsb.gov/speeches/former/hall/jhc000823.htm.

[615] National Transportation Safety Board. Safety Recommendations Issued by Mode 1967 to September 1, 2001. Technical report, NTSB, Washington, DC United States of America, 2001. http://www.ntsb.gov/Recs/statistics/safety_recommendations_issued_by.htm.

[616] National Transportation Safety Board. Status of Safety Recommendations by Mode of Transportation With Acceptance Rates as of September 1, 2001. Technical report, NTSB, Washington, DC United States of America, 2001. http://www.ntsb.gov/Recs/statistics/acceptance_rate_modal.htm.

[617] National Transportation Safety Board. About the NTSB: The Investigative Process. Technical report, NTSB, Washington, DC United States of America, 2002. http://www.ntsb.gov/abt%5Fntsb/invest.htm.

[618] National Transportation Safety Board. Fire on board the U.S. Passenger Ferry Columbia, Chatham Strait Near Juneau, Alaska June 6, 2000. Technical Report Marine Accident Report PB2001-916403, NTSB/MAR-01/02, NTSB, Washington, DC United States of America, 2002. http://www.ntsb.gov/Publictn/2001/MAR0102.pdf.

[619] National Transportation Safety Board of Canada. A safety study of the operational relationship between ship masters/ watchkeeping officers and marine pilots. Technical Report SM9501, NTSB Canada, Hull, Quebec, Canada, 1995. http://www.bst.gc.ca/ENG/reports/marine/ems9701.html.

[620] National Transportation Safety Board of Canada. Aviation Occurrence Report: Operating Irregularity Air BC Ltd. British Aerospace BAE146 Flight ABL814 Vancouver, British Columbia 30 nm N 01 February 1998. Technical Report A98P0018, NTSB Canada, Hull, Quebec, Canada, 1998. http://www.bst.gc.ca/eng/reports/air/1998/ea98p0018.html.

[621] National Transportation Safety Board of Canada. Marine Investigation Report: Engine-Room Fire The Self-unloading Bulk Carrier "NANTICOKE" At 39 degrees 20' N, 072 degrees 22' W Western North Atlantic Ocean 20 July 1999. Technical Report M99F0023, NTSB Canada, Hull, Quebec, Canada, 1999. http://www.bst.gc.ca/ENG/reports/marine/1999/m99f0023/em99f0023.htm.

[622] National Transportation Safety Board of Canada. The Securitas confidential incident reporting system. Technical report, NTSB Canada, Hull, Quebec, Canada, 2000. http://www.bst.gc.ca/ENG/.

[623] National Transportation Safety Board of Canada. Score your safety culture. Technical report, NTSB Canada, Hull, Quebec, Canada, 2002. http://www.tc.gc.ca/civilaviation/systemsafety/tp13844/menu.htm.

[624] NATO, Department of Peacekeeping Operations. The Comprehensive Report on Lessons Learned from United Nations Operation in Somalia (UNOSOM). Technical report, North Atlantic Treaty Organisation, New York, USA, 1995. http://www.un.org/Depts/dpko/lessons/mandate.htm.

[625] NATO, Department of Peacekeeping Operations. Comprehensive Report on Lessons Learned from United Nations Assistance Mission for Rwanda (UNAMIR). Technical report, North Atlantic Treaty Organisation, New York, USA, 1996. http://www.un.org/Depts/dpko/lessons/rwandisc.htm.

[626] J. Von Neumann and O. Morgenstern. *The Theory Of Games And Economic Behaviour.* Princeton University Press, Princeton, United States of America, 1944.

[627] P.G. Neumann. *Computer Related Risks.* Addison Wesley, Reading, MA, United States of America, 1995.

[628] New Zealand Transport Accident Investigation Commission. Annual Report 2000-2001. Technical report, TAIC, Wellington, New Zealand, 2000. http://www.taic.org.nz/reports/ar_2000-01-report.html.

[629] New Zealand Transport Accident Investigation Commission. Marine Safety Recommendations and Responses: Report 00-211, harbour tug Waka Kume, loss of control, Auckland Harbour, 19 November 2000. Technical report, TAIC, Wellington, New Zealand, 2000. http://www.taic.org.nz/marine/recommend_00-211.html.

[630] New Zealand Transport Accident Investigation Commission. Marine safety recommendations and replies since January 1997. Technical report, TAIC, Wellington, New Zealand, 2001. http://www.taic.org.nz/marine/recommendations-m.html.

[631] New Zealand Transport Accident Investigation Commission. Simplified flow chart of report and safety recommendation progress. Technical report, TAIC, Wellington, New Zealand, 2001. http://www.taic.org.nz/aboutus/flowchart.html.

[632] New Zealand Transport Accident Investigation Commission. How to obtain a TAIC occurrence report. Technical report, TAIC, Wellington, New Zealand, 2002. http://www.taic.org.nz/aboutus/obtaining.html.

[633] NHS Expert Group on Learning from Adverse Events in the NHS. An organisation with a memory. Technical report, National Health Service, London, United Kingdom, 2000. www.doh.gov.uk/orgmemreport/index.htm.

[634] J. Nielsen. *Usability Engineering.* Academic Press, San Diego, CA, United States of America, 1993.

[635] I. Nonaka and H. Takeuchi. *The Knowledge Creating Company.* Oxford University Press, Oxford, UK, 1995.

[636] D.A. Norman. The ‘problem’ with automation : Inappropriate feedback and interaction not ‘over-automation’. In D.E. Broadbent, J. Reason, and A. Baddeley, editors, *Human Factors In Hazardous Situations*, pages 137–145. Clarendon Press, Oxford, United Kingdom, 1990.

[637] B.E. Norris, D.Z. Rashid, and B.L.W. Wong. Wayfinding/navigation within a QTVR virtual environment. In M. A. Sasse and C.W. Johnson, editors, *Interact ’99*, pages 544–551, Amsterdam, NL, 1999. IOS Press.

[638] Y.I. Noy. *Ergonomics and Safety of Intelligent Driver Interfaces.* Lawrence Erlbaum, Mahway, NJ, USA, 1997.

[639] H. Mu noz Avila, J.A. Hendler, and D.W. Aha. Conversational case-based planning. *Review of Applied Expert Systems*, 5:163–174, 1999.

[640] Nuclear Installations Inspectorate. HSE Team Inspection of the Control and Supervision of Operations at BNFL’s Sellafield Site. Technical report, Health and Safety Executive, Nuclear Safety Directorate Information Centre, Bootle, United Kingdom, 2000. http://www.hse.gov.uk/nsd/team.htm.

[641] Nuclear Installations Inspectorate. An investigation into the falsification of pellet diameter data in the MOX Demonstration Facility at the BNFL Sellafield site and the effect of this on the safety of MOX fuel in use. Technical report, Health and Safety Executive, Nuclear Safety Directorate Information Centre, Bootle, United Kingdom, 2000. http://www.hse.gov.uk/nsd/mox1.htm.

[642] Occupational Safety and Health Administration. Secretary of Labor (Complainant) vs Gearhart-Owen Industries, Inc. (Respondent). Technical Report OSHRC Docket No. 4263, Case Citation 10 BNA OSHC 2193, 192 CCH OSHD Paragrah 26,329, US Department of Labour, Washington DC, United States of America, 1982. http://www.osha-slc.gov/REVIEW_data/D19821122.html.

[643] Occupational Safety and Health Administration. Secretary of labor (complainant) versus all purpose crane inc. (respondent). Technical Report OSHRC Docket 82-284, Case Citation: 13 BNA OSHC 1236, US Department of Labour, Washington DC, United States of America, 1987. http://www.osha-slc.gov/REVIEW_data/D19870414.html.

[644] Occupational Safety and Health Administration. Secretary of labor (complainant) versus keco industries inc. (respondent). Technical Report OSHRC Docket No. 81-263, Case Citation: 13 BNA OSHC 1161, US Department of Labour, Washington DC, United States of America, 1987. http://www.osha-slc.gov/REVIEW_data/D19870327.html.

[645] Occupational Safety and Health Administration. Secretary of Labor (Complainant) vs All Purpose Crane Inc. (Respondent). Technical Report OSHRC Docket No. 82-284, Case Citation 13 BNA OSHC 1236, US Department of Labour, Washington DC, United States of America, 1987. http://www.osha-slc.gov/REVIEW_data/D19870414.html.

[646] Occupational Safety and Health Administration. Secretary of labor (complainant) vs zunker contractors inc. (respondent). Technical Report OSHRC Docket No. 82-936, Case Citation: 13 BNA OSHC 2200, US Department of Labour, Washington DC, United States of America, 1989. http://www.osha-slc.gov/REVIEW_data/D19890427.html.

[647] Occupational Safety and Health Administration. OSHA response to significant events of potentially catastrophic consequences. Technical Report OSHA Directives CPL 2.94, US Department of Labour, Washington DC, United States of America, 1991. http://www.osha-slc.gov/OshDoc/Directive_data/CPL_2_94.html.

[648] Occupational Safety and Health Administration. OSHA high injury/illness rate targeting and cooperative compliance programs. Technical Report OSHA Directives, CPL 2-0.119, US Department of Labour, Washington DC, United States of America, 1997. http://www.osha-slc.gov/OshDoc/Directive_data/CPL_2-0_119.html.

[649] Occupational Safety and Health Administration. Osha's small business outreach training program instructional guide. Technical report, US Department of Labour, Washington DC, United States of America, 1997. http://www.osha-slc.gov/SLTC/smallbusiness/sec6.html.

[650] Occupational Safety and Health Administration. Occupational exposure to 1,3-butadiene. Technical Report 29 Code of Federal Regulations Parts 1910, 1915 and 1926, [Docket No. H-041], US Department of Labour, Washington DC, United States of America, 1999. http://www.osha-slc.gov/Preamble/13Butadiene_data/1_3_BUTADIENE5.html.

[651] Occupational Safety and Health Administration. OSHA's proposed ergonomics standard. Technical report, US Department of Labour, Washington DC, United States of America, 1999. http://www.osha-slc.gov/ergonomics-standard/fs-over.html.

[652] Occupational Safety and Health Administration. New ways of working: OSHA facts. Technical report, US Department of Labour, Washington DC, United States of America, 2000. http://www.osha-slc.gov/OSHAFacts/OSHAFacts2_3_00.pdf.

[653] Occupational Safety and Health Administration. OSHA record keeping guidelines. Technical report, US Department of Labour, Washington DC, United States of America, 2000. http://www.osha-slc.gov/rkeep_data/RKEEP_10.html.

[654] Occupational Safety and Health Division. Attention all health and safety reps: beginner's guide to effective accident investigation. Technical report, Department of Labour, Wellington, New Zealand, 1989. http://www.nohsc.gov.au/publications/pamphlets/a/002146.htm.

[655] Office of Occupational Safety and Health Policy. Computerized accident/incident reporting system. Technical report, US Department of Energy, 2000. http://tis.eh.doe.gov/cairs/cairs/facts.htm.

[656] Office of System Safety. The global aviation information network (GAIN). Technical report, U. S. Federal Aviation Administration, Washington DC, United States of America, 2002. http://www.asy.faa.gov/gain.

[657] K. O'Hara and A. Sellen. A comparison of reading paper and on-line documents papers: Papers about paper. In *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems*, pages 335–342, New York, USA, 1997. ACM Press.

[658] M. O'Leary. The British Airways human factors reporting programme. In V. de Keyser and D. Javaux, editors, *Proceedings of the Third Workshop on Human Error, Safety and Systems Development*, Liège, Belgium, 1999. ULg, University of Liège.

[659] M. O'Leary. Situation awareness and automation. In K. Abbott, J.-J. Speyer, and G.Boy, editors, *HCI Aero 2000: International Conference on Human-Computer Interfaces in Aeronautics*, pages 73–79, Toulouse, France, 2000. Cepadues-Editions.

[660] M. O'Leary and S.L. Chappell. Confidential incident reporting systems create vital awareness of safety problems. *ICAO Journal*, pages 11–13, 1996.

[661] J. Orsanu. Shared mental models and crew decision making. Technical Report 46, Cognitive Sciences Laboratory, Princeton University, Princeton, NJ, United States of America, 1990.

[662] J. Orsanu, U. Fischer, L.K. McDonnell, J. Davison, K. Haars, E. Villeda, and C. Van Aken. Detecting and correcting errors in flight. Technical report, NASA Ames Research Centre, Human Factors Group, Moffett Field, CA, United States of America, 1999. http://olias.arc.nasa.gov/reports/detecting.html.

[663] P. Palanque and R. Bastidee. Synergistic modelling of tasks, system and users using formal specification techniques. *Interacting With Computers*, 9(12):129–153, 1997.

[664] A. Pasquini, A. Rizzo, and L. Save. Analysis of incidents involving interactive systems. Technical report, ENEA, Rome, Italy, 2000.

[665] M. E. Pate-Cornell. Learning from the Piper Alpha accident : A postmortem analysis of technical and organisation factors. *Risk Analysis*, 13(2):215–232, 1993.

[666] F. Paterno. *Model-Based Design and Evaluation of Interactive Applications.* Springer Verlag, Berlin, Germany, 2000.

[667] M. Patten. More than meets the eye. Technical Report 7, Aviation Safety Reporting System, NASA Ames Research Centre, California, United States of America, September 1995. http://asrs.arc.nasa.gov./directline_issues/dl7_laser.htm.

[668] R.D. Patterson. Auditory warning sounds in the work environment. In D.E. Broadbent, J. Reason, and A. Baddeley, editors, *Human Factors In Hazardous Situations*, pages 37–44. Clarendon Press, Oxford, United Kingdom, 1990.

[669] B.J. Payne. Dealing with hazard and risk in planning. In R.F. Griffiths, editor, *Dealing with Risk.* Manchester University Press, Manchester, UK, 1981.

[670] D. G. Payne, M. J. Wenger, and M. S. Cohen. Cognitive processing and hypermedia comprehension: A preliminary synthesis. In *Proceedings of the Fifth International Conference on Human-Computer Interaction*, volume 2, pages 633–638, 1993.

[671] J. Pearl. *Causality: Models, Reasoning and Inference.* Cambridge University Press, Cambridge, United Kingdom, 2000.

[672] M. Pedralli. *Vers un Environnement Multimedia Pour L'Analyse Video des Causes d'Erreurs Humaines Application dans les Simulateurs d'Avions.* PhD thesis, LIHS, University of Toulouse 1, Toulouse, France, 1996.

[673] C. Perin. Organisations as contexts: Implications for safety science and practice. *Industrial and Environmental Crisis Quarterly*, 9(2):152–174, 1995.

[674] W.D. Perreault and L.E. Leigh. Reliability of normal data based on qualitative judgements. *Journal of Marketing Research*, pages 135–148, 1989.

[675] C. Perrow. *Normal Accidents: Living with High-Risk Technologies.* Princeton University Press, Princeton, NJ, United States of America, 1999.

[676] Peter Majgrd Nrbjerg. The creation of an aviation safety reporting culture in danish air traffic control. In C.M.Holloway C.J. Hayhurst and B. Strauch, editors, *2nd Workshop on the Investigation and Reporting of Incidents and Accidents 2003*, NASA/CP-2003-212642, pages 153–164. NASA Langley Research Centre, Virginia, USA, 2003.

[677] J. Petersen. Focus and causal reasoning in disturbance management of complex, dynamic systems. In P.C. Cacciabue, editor, *Human Decision Making and Manual Control: EAM2000*, EUR 19599 EN, pages 43–49. European Commission, Joint Research Centre, Ispra, Italy, 2000.

[678] J.L. Peterson. Petri nets. *Computing Surveys*, 9(3):223 – 252, 1977.

[679] H. Petroski. *Design Paradigms: Case Histories of Error and Judgement in Engineering.* Cambridge University Press, Cambridge, UK, 1994.

[680] E.H. Phillips. GAIN transitions to airline operations. *Aviation Week and Space Technology*, August 7 2000. http://www.AviationNow.com.

[681] N. Pidgeon and M. O'Leary. Organisational safety culture: Implications for aviation practice. In *Aviation Psychology in Practice*. Ashgate Publishing, Aldershot, UK, 1994.

[682] D.A. Pietro, L.J. Shyavitz, R.A. Smith, and B.S. Auerbach. Detecting and reporting medical errors: why the dilemma? *British Medical Journal*, 320(7237):794–796, 2000.

[683] S. J. Popkin, V. E. Gwiasda, J. M. Amendolia, A. A. Anderson, W. A. Johnson G. Hanson, E. Martel, L. M. Olson, and D. P. Rosenbaum. The hidden war: The battle to control crime in Chicago's public housing. *National Institute of Justice Journal*, page 19, March 1998. http://www.ncjrs.org/pdffiles/jr000235.pdf.

[684] T.R. La Porte and P.M. Consolini. Working in practice but not in theory: Theoretical challenges of high reliability organisations. *Journal of Public Administration Research and Theory*, 1(1):19–47, 1991.

[685] M.I. Posner, J.M. Nissen, and R. Klein. Visual dominance: an information processing account of its origins and significance. *Psychological Review*, pages 157–171, 1976.

[686] J. Preece, Y. Rogers, H. Sharp, D. Benyon, S. Holland, and T. Carey. *Human Computer Interaction*. Addison Wesley, Wokingham, UK, 1994.

[687] R. Price. The seafood network at UC DAVIS. Technical report, Food Science and Technology Dept, University of California, Davis, CA, USA, 2001. http://vm.cfsan.fda.gov/ frf/seanetls.html.

[688] Prime Minister's Press Office. Fact Files: Improving Rail Safety. Technical report, 10 Downing Street, London, UK, 2001. http://www.number-10.gov.uk/default.asp?pageid=2679.

[689] C. Puppe. *Distorted Probabilities And Choice Under Risk*. Lecture Notes In Economics And Mathematical Systems, No 363. Springer Verlag, Berlin, Germany, 1991.

[690] Railtrack. The Ladbroke Grove Rail Enquiry: Fact Sheets. Technical report, Railtrack, 1999. http://www.railtrack.co.uk/cullen/index.html.

[691] Railway Safety, (Controller, Safety Strategy & Risk). Railway Group 2001/02 Half-year Safety Performance Report. Technical report, Railway Group, London, UK, 2001. http://www.railwaysafety.org.uk/q2spr.asp.

[692] A. Ram. Indexing, elaboration and refinement: Incremental learning of explanatory cases. In J.L. Kolodner, editor, *Case-based Learning*. Kluwer, London, 1993.

[693] R. Randell. Coping strategies in an intensive care unit. *Cognition, Technology and Work*, 5(3):163–170, 2003.

[694] J. Rasmussen. Skills, rules and knowledge; signals, signs and symbols and other distinctions in human performance models. *IEEE Transactions On Systems, Man And Cybernetics*, SMC-13(3):257–266, 1983.

[695] J. Rasmussen. *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. Elsevier, North Holland, 1986.

[696] J. Rasmussen. Diagnostic reasoning in action. *IEEE Transactions on Systems, Man and Cybernetics*, 23(4):981–992, 1993.

[697] J.O. Rawlings, S.G. Pantula, and D.A. Dickey. *Applied Regression Analysis: A Research Tool*. Springer Texts in Statistics, London, UK, 1998.

[698] J. Reason. The contribution of latent human failures to the breakdown of complex systems. *Philosophical Transactions of the Royal Society of London*, pages 475–484, 1990.

[699] J. Reason. *Human Error*. Cambridge University Press, Cambridge, UK, 1990.

[700] J. Reason. Foreword. In M.S. Bogner, editor, *Human Error in Medicine*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA, 1994.

[701] J. Reason. *Managing the Risks of Organizational Accidents*. Ashgate Publishing, Aldershot, UK, 1997.

[702] P. E. Reimers and S. M. Chung. Intelligent user interface for very large relational databases. In M. J. Smith and G. Salvendy, editors, *Proc. of the Fifth International Conference on Human-Computer Interaction*, volume 2, pages 134–139, Holland, 1993. Elsevier Science.

[703] R. Remington, M. Shafto, and M. Freed. Making human-machine system simulation a practical engineering technique. Technical report, Human Factors Research and Technology Division, NASA Ames Research Center, Moffett Field, USA, 2000. http://human-factors.arc.nasa.gov:80/reports/makinghuman.html.

[704] Research and Special Programs Administration. Transportation Department Takes Action to Improve Pipeline Accident Reporting. Technical report, Office of Public Affairs, Department of Transportation, Washington DC, United States of America, 2001. http://www.dot.gov/affairs/rspa901.htm.

[705] I.J. Rimson and L. Benner. Quality controls for the investigation/prevention process. In *Symposium on High Consequence Operations Safety*, Albuquerque, New Mexicom USA, 1994. http://www.iprr.org/Papers/Sandia/SANDIA4.html.

[706] K. Risden, M. Czerwinski, T. Munzner, and D. Cook. An initial examination of ease of use for 2D and 3D information visualizations of web content. *International Journal of Human-Computer Studies*, 53(5):695–714, 2000. Special Issue on Empirical Evaluations of Information Visualizations.

[707] J. Robb. Annual report 1997-98. Technical report, British Energy, London, United Kingdom, 1998. http://www.british-energy.co.uk/environment/mn_hsed_safety.html.

[708] F.W. Robbertze. The winning edge: Integrity and excellence in the Department of Defence (DOD). Technical report, South African Defence College, Thaba Tshwane, Pretoria, South Africa, 1999. http://www.mil.za/.

[709] Roben. Safety and health at work: Report of the Roben committee 1970-1972. Technical report, Her Majesty's Stationary Office, London, United Kingdom, 1972.

[710] K.H. Roberts. New challenges in organizational research: High reliability organizations. *Industrial Crisis Quarterly*, 3(2):111–125, 1989.

[711] T. H. Rockwell. Some exploratory research on risk acceptance in a man-machine setting. *Journal of the American Society of Safety Engineers*, pages 23–29, 1962.

[712] M. Rodgers and D. Duke. SATORI: Situation assessment through the recreation of incidents. Technical Report DOT/FAA/AM-93/12, US Federal Aviation Administration, Office of Aviation Medicine, Washington, DC, United States of America, 1993.

[713] W.P. Rogers. Report of the Presidential commission on the space shuttle Challenge accident. Technical Report Executive Order 12546 of February 3, 1986, US Government Accounting Office, Washington, DC, USA, 1986. http://science.ksc.nasa.gov/shuttle/missions/51-l/docs/rogers-commission/.

[714] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiations Among Computers*. MIT Press, Cambridge, United States of America, 1994.

[715] Royal College of Anaesthetists. Critical incident form. Technical report, Professional Standards Directorate, The Royal College of Anaesthetists, London, UK, January 1998. http://www.rcoa.ac.uk/critincident/citext.html.

[716] B. Russell. On the notion of cause. *Proceedings Of The Aristotelean Society*, 13:1–25, 1912.

[717] Safety Regulatory Commission. Safety regulatory requirement: Reporting and assessment of safety occurrences in ATM. Technical report, EUROCONTROL, Brussels, Belgium, 1999. http://www.eurocontrol.be/dgs/src/en/index.html.

[718] S.D. Sagan. *The Limits of Safety: Organisations, Accidents and Nuclear Weapons.* Princeton University Press, Princeton, NJ, United States of America, 1993.

[719] M. Sage and C.W. Johnson. A declarative prototyping environment for multi-user, safety-critical systems. In J. Dixon, editor, *Proceedings of the 17th International Systems Safety Conference*, pages 614–623, Unionville, Virginia, United States of America, 1999. The Systems Safety Society.

[720] M. Sage and C.W. Johnson. Formally verified rapid prototyping for air traffic control. In V. de Keyser and D. Javaux, editors, *Proceedings of the Third Workshop on Human Error, Safety and Systems Development*, Liège, Belgium, 1999. Psychologie du Travail et des Entreprises, Université de Liège.

[721] J.G. Salaman. *Community and Occupations: An Exploration of Community and Work Leisure Realtionships.* Cambridge Papers in Sociology Number 4. Cambridge University Press, Cambridge, 1974.

[722] S. Salminen, T. Klen, and K. Ojanen. Risk-taking behaviour of forestry workers. In S.A. Robertson, editor, *Contemporary Ergonomics*, pages 401–405. Taylor and Francis, London, United Kingdom, 1997.

[723] San Jacinto Rail Limited. Transportation safety. Technical report, San Jacinto Rail Limited Partners, Houston, Texas, USA, 2000. http://www.sanjacintorail.com/safety.html.

[724] P.M. Sanderson and C. Fisher. Exploratory sequential data analysis: Foundations. *Human-Computer Interaction*, 9:251–317, 1994.

[725] W.E. Saris. *Computer Assisted Interviewing.* Sage, Newbury Park, USA, 1991.

[726] N.B. Sarter and D.D. Woods. Situation awareness: A critical but ill-defined phenomenon. *International Journal of Aviation Psychology*, 1:43–55, 1991.

[727] N.B. Sarter and D.D. Woods. Teamplay with a powerful and independent agent: A full-mission simulation study. *Human Factors*, 42(3):390–402, 2001.

[728] K. Sasou and J. Reason. Team errors: Definition and taxonomy. *Reliability Engineering and System Safety*, pages 1–9, 1999.

[729] M. Schiavo. *Flying Blind, Flying Safe.* Avon, New York, United States of America, 1997.

[730] D. Schofield, J. Noond, L. Goodwin, K. Fowle, and M. Doyle. How real is your reconstruction? Developments in computer graphics and virtual reality. In *Forensic techniques for the 21st Century*, London, UK, 2000. Lloyds of London.

[731] B. Scott and B. Nelson. Release of board of inquiry findings and action into the death of an army cadet. Technical report, Department of Defence, Canberra, Australia, 2001. Min 096/01, http://www.defence.gov.au/minister/2001/96180401.doc.

[732] A. Scrivener. Special report on Ladbroke Grove: ‘Pass the signal - pass the blame’. *The Locomotive Journal*, pages 8–9, June 2000. Quoted extracts from evidence to Lord Cullen’s inquiry into the Ladbroke Grove accident, http://www.aslef.org.uk/dox/loco_june_00.pdf.

[733] T. Seamster. Automation and advanced crew resource management. In S. Dekker and E. Hollnagel, editors, *Coping with Computers in the Cockpit*, pages 195–213. Ashgate, Brookfield, VM, United States of America, 1999.

[734] T.L. Seamster, D.A. Boehm-Davis, R.W. Holt, and K. Schultz. Developing advanced crew resource management (ACRM) training: A training manual. Technical Report Human Factors AAR-100, Federal Aviation Administration, Washington DC, United States of America, 1998. http://www.hf.faa.gov/products/dacrmt/DACRMT.pdf.

[735] J.B. Sexton, E.J. Thomas, and R.L. Helmreich. Error, stress, and teamwork in medicine and aviation: cross sectional surveys. *British Medical Journal*, 320(7237):745–749, 2000. http://www.bmj.com/cgi/content/full/320/7237/745.

[736] Sheen. Report of Court Number 8074: MV Herald of Free Enterprise. Technical report, Department of Transport, Her Majesty's Stationery Office, London, United Kingdom, 1987.

[737] T.B. Sheridan. On how often the supervisor should sample. *IEEE Transactions of Systems, Sciences and Cybernetics*, SSC-6:140–145, 1972.

[738] T.B. Sheridan. Understanding human error and aiding human diagnostic behaviour in nuclear power plants. In J. Rasmussen and W.B. Rouse, editors, *Human Detection and Diagnosis of System Failures*. Plenum Press, New York, United States of America, 1981.

[739] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualisation. In *IEEE Conference on Visual Languages*, pages 336–343, 1996.

[740] B. Shneiderman. *Desigining the User Interface*. Addison Wesley, Reading, MA, USA, 1998.

[741] Singapore Army. Preventing heat injuries: The commanders guide. Technical report, Safety Organisation, General Staff Inspectorate, Singapore, 2001. http://www.mindef.gov.sg/army/gsi/education/heat/heatprevent.pdf.

[742] Singaporean Ministry of Manpower. How to report accidents to the Ministry of Manpower. Technical report, Ministry of Manpower, Singapore, 2000. http://www.gov.sg/mom/wpeaw/isfile/is3a.htm.

[743] A. Singhal, G. Salton, M. Mitra, and C. Buckley. Document length normalization. *Information Processing and Management*, 32(5):619–633, 1996.

[744] G. Sitwell and S. Purcel. Assessment of Investigations into Signals Passed at Danger (SPADs). Technical Report BL2077 004 TR06, WS Atkins Rail Limited, under contract from the HSE, London, UK, 2001. http://www.hse.gov.uk/railway/spad/spadrep1.pdf.

[745] W. Smith, J. Dowell, and M. Ortega-Lafuente. Narratran: a tool for designing emergency management training. In T.R. G. Green, L. Bannon, C. Warren, and J. Buckley, editors, *Cognition and Cooperation: Proceedings of the Ninth European Conference on Cognitive Ergonomics*, pages 97–102. European Association for Cognitive Ergonomics (EACE), Paris, 1998. ISBN-874653-48-8.

[746] G. Smithsimon. Ethnograffiti: Graffiti, writers, and space at the Phun Phactory. Technical report, Department of Sociology, Columbia State University, USA, 1997. http://www.columbia.edu/∼gs228/writing/ethnograffiti.htm.

[747] S. A. Snook. *Friendly Fire*. Princeton University Press, 2000.

[748] P. Snowdon. *The Use of Argumentation and Rhetoric in Accident Reports*. PhD thesis, Department of Computing Science, University of Glasgow, Glasgow, Scotland, 2002.

[749] P. Snowdon and C.W. Johnson. Results of a preliminary survey into the usability of accident and incident reports. In J. Noyes and M. Bransby, editors, *People in Control: An international conference on human interfaces in control rooms, cockpits and command centres*, pages 258–262, Savoy Place, London, United Kingdom, 1999. The Institute of Electrical Engineers. Bath, UK, 21-23 June 1999.

[750] Society of Automotive Engineers. Accident reconstruction: Technology and animation. Technical report, SAE Publishing, Washington, DC, USA, 1999. ISBN 0768003393.

[751] I. Sommerville, T. Rodden, P. Sawyer, and R. Bentley. Sociologists can be surprisingly useful in interactive systems design cooperative systems. In A. Monk, D. Diaper, and M.D. Harrison, editors, *HCI'92: People and Computers VII*, pages 341–353, Cambridge, UK, 1992. Cambridge University Press.

[752] F. Speirs and C.W. Johnson. Designing information visualisations for incident databases. In J.H. Wiggins and S.G. Thompson, editors, *Proceedings of the 20th Annual Systems Safety Conference*, pages 280–290, Unionville, VA, United States of America, 2002. Systems Safety Society.

[753] F. Speirs, B. Robinson, and C.W. Johnson. The validation of information visualisation in incident reporting systems. In C.W. Johnson, editor, *Proceedings of the First Workshop on the Investigation and Analysis of Incidents and Accidents (IRIA-2002)*, Department of Computing Science, University of Glasgow, Scotland, 2002. Glasgow Accident Analysis Group. Accepted and to appear.

[754] R. F. Spohrer and D. M. Maciejewski. Modus operandi for the expert witness. *Chemical Health and Safety*, 3(1):8–12, 1996. http://pubs.acs.org/hotartcl/chas/96/janfeb/janfeb.html.

[755] S. Staender. Critical Incidents Reporting System (CIRS): Critical incidents in anaesthesiology. Technical report, Department of Anaesthesia, University of Basel, Basel, Switzerland, 2000. http://www.medana.unibas.ch/cirs/.

[756] S. Staender, M. Kaufman, and D. Scheidegger. Critical incident reporting in anaesthesiology in switzerland using standard internet technology. In C.W. Johnson, editor, *1st Workshop on Human Error and Clinical Systems*, Glasgow, Scotland, 1999. Department of Computing Science, University of Glasgow. http://www.dcs.gla.ac.uk/ johnson/papers/HECS_99/Standaer_Kaufman_Scheidegger.htm.

[757] D. Stanyer. *Redefining the Hyperlink*. PhD thesis, Division of Informatics, Edinburgh, UK, 2001.

[758] N.M. Steblay. A meta-analytic review of the weapon foccus effect. *Law and Human Behaviour*, pages 413–424, 1992.

[759] N.M. Steblay. Social influence in eyewitness recall: A meta-analytic review of lineup instruction effects. *Law and Human Behaviour*, pages 283–298, 1997.

[760] E.S. Stein and B. Rosenberg. The measurement of pilot workload. Technical Report Report DOT/FAA/CT82-23, NTIS No. ADA124582, Federal Aviation Authority, Atlantic City, United States of America, 1983.

[761] R.B. Stillman, J.B. Stephenson, K.A. Rhodes, K.J. Daubenspeck, P.R. Dugan, C.T. Chaplain, R.P. Kissel, and F. Maguire. Information security: Computer attacks at department of defense pose increasing risks. Technical Report AIMD-96-84, Accounting and Information Management Division, United States General Accounting Office and Intelligence Resource Program, Washington, DC, USA, 1996. http://www.fas.org/irp/gao/aim96084.htm.

[762] N. Storey. *Safety-Critical Computer Systems*. Addison Wesley, Reading MA, United States of America, 1996.

[763] B. Strauch. Correspondence with Peter Ladkin, 1997. Cited in [470].

[764] J. Suakas. The role of management in accident prevention. In *First International Congress on Industrial Engineering and Management*, Paris, France, 11-13 June 1986. Cited in [486].

[765] P. Suppes. *A Probabilistic Theory of Causality.* Elsevier, Amsterdam, North Holland, 1970.

[766] A.D. Swain and H.E. Guttman. Handbook of human reliability with emphasis on nuclear power plant applications. Technical Report NUREG CR-1278, US Nuclear Regulatory Commission, Washington DC, United States of America, 1983.

[767] Swedish Board of Accident Investigation. Collision between the Swedish Vessel MT Tärnsjö and the Russian Vessel MV Amur-2524, on 6th February 1998 at Strömskär C County Sweden, S01/98. Technical Report S 1998:02e, Statens Haverikommission, Stockholm, Sweden, 1998. http://www.havkom.se/rapportSammandrag/s1998e02.pdf.

[768] Swedish Board of Accident Investigation. Statens Haverikommission - the board of accident investigation. Technical report, Statens Haverikommission, Stockholm, Sweden, 2001. http://www.havkom.se/english.phtml.

[769] S.Welch and P.A. Rubini. Three dimensional simulation of a fire resistance furnace. In *Proceedings of 5th International Symposium on Fire Safety Science*, Melbourne, Australia, 1987. International Association for Fire Safety Science.

[770] A.J. Tattersall. Individual differences in performance. In M.W. Smolensky and E.S. Stein, editors, *Human Factors in Air Traffic Control*, pages 185–213. Academic Press, London, United Kingdom, 1998.

[771] J.D. Teasdale and P.J. Barnard. *Affect, Cognition and Change.* Lawrence Erlbaum Associates, London, United Kingdom, 1993.

[772] Thames Regional Health Authority. Report of the inquiry into the London Ambulance Service Computer Aided Dispatch System. Technical report, Thames Regional Health Authority, Communications Directorate, South West Thames Regional Health Authority, February 1993.

[773] B. Toft and S. Reynolds. *Learning from Disaster: A Management Approach.* Perpetuity Press, Leicester, UK, 1997.

[774] C. Tomlin. *Hybrid Control of Air Traffic Management Systems.* PhD thesis, Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkley, United States of America, 1998.

[775] S. Toulmin. *The Uses of Argument.* Cambridge University Press, Cambridge, United Kingdom, 1958.

[776] Transport Canada. Hot bearing detector study. Technical Report TP 12691-E, Railway Safety Directorate, Quebec, Canada, 1996. http://www.tc.gc.ca/railway/ENGLISH/HOT_BEARING/hbdrp6_1.htm.

[777] Transport Canada. Uncontrolled movements of railway equipment. Technical Report Railway Safety Facts 1996, Railway Safety Directorate, Quebec, Canada, 1996. http://www.tc.gc.ca/railway/ENGLISH/RUNAWAY/Runaway_1.htm.

[778] Transport Canada. Departmental performance report for the period ending march 31, 2001. Technical report, TC, Quebec, Canada, 2001. http://www.tc.gc.ca/estimate/dpr/00-01/en/dpr_printableversion_e.htm.

[779] Transport Canada. Railway Safety Act: Railway safety management system regulations - interpretation. Technical report, TC, Quebec, Canada, 2001. http://www.tc.gc.ca/railway/SMS_Regulations.htm.

[780] Transport Canada. Railway safety management system guide. Technical report, TC, Quebec, Canada, 2001. http://www.tc.gc.ca/railway/publications/smsguide.htm.

[781] Transport Salaried Staffs' Association. TSSA Reps Bulletin June 2001: Accidents @ Work? Technical report, TSSA, London, U.K., 2001. http://www.tssa.org.uk/advice/hs/hs12.htm.

[782] Transport South Australia. Rail Safety Act 1996, Notification of Occurrences, Incident Normalising Factors Forms. Technical report, Rail Safety Unit, Government of Southern Australia, Walkerville, Australia, 1999. http://www.transport.sa.gov.au/rail/normfact.htm.

[783] Transportation Safety Board of Canada. Railway Safety Facts 1996. Technical report, Safety and Security, Rail Safety Directorate, Rail Safety Programs Branch, TSB, Hull, Quebec, Canada, 1996. http://www.tc.gc.ca/railway/ENGLISH/DERAIL/DERAIL_3.HTM.

[784] Transportation Safety Board of Canada. Overturning of the Pilot Boat "NAVIMAR V" Port of Québec, Quebec 07 August 1997. Technical Report M97L0076, TSB, Hull, Quebec, Canada, 1997. http://www.bst.gc.ca/ENG/reports/marine/1997/m97l0076/em97l0076.htm.

[785] Transportation Safety Board of Canada. Bottom Contact Bulk Carrier "ALAM SELAMAT" Fraser River, British Columbia 16 June 1999. Technical Report M99W0087, TSB, Hull, Quebec, Canada, 1999. http://www.tsb.gc.ca/ENG/reports/marine/1999/m99w0087/em99w0087.htm.

[786] Transportation Safety Board of Canada. Striking of a Dock with an Unloading Boom, the Self-unloading Bulk Carrier 'ALGOBAY', Poe Lock, Sault Ste. Technical Report M99F0042, TSB, Hull, Quebec, Canada, 1999. http://www.tsb.gc.ca/ENG/reports/marine/1999/m99f0042/em99f0042.htm.

[787] Transportation Safety Board of Canada. Annual statistical summary of railway occurrences (2000). Technical report, TSB, Hull, Quebec, Canada, 2001. http://www.tsb.gc.ca/eng/stats/rail/year00/rail-eng.htm.

[788] Transportation Safety Board of Canada. New released reports. Technical report, TSB, Hull, Quebec, Canada, 2001. http://www.tsb.gc.ca/ENG/reports/marine/RptMar_Indx_New.html.

[789] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2001.

[790] B.A. Turner. The sociology of safety. In D.I. Bockley, editor, *Engineering Safety*, pages 186–201. McGraw Hill, London, United Kingdom, 1992.

[791] H.R. Turtle and W.B. Croft. Evaluation of an inference network-based retrieval model. *ACM Transactions on Information Systems*, 9(3):187–222, 1991.

[792] UK House of Commons Select Committee on Defence. Government observations on the fourteenth report from the committee, session 1999-2000, on the lessons of Kosovo. Technical Report HC 347-I, UK House of Commons, London, UK, 2000. http://www.parliament.the-stationery-office.co.uk/pa/cm200001/cmselect/cmdfence/178/17805.htm.

[793] UK Marine Pollution Control Unit. The Sea Empress Incident: Summary of Report. Technical report, UK Coastguard Agency and the Department of the Environment, Transport and the Regions, Southampton, UK, 1997. http://www.shipping.dtlr.gov.uk/tca/sea_emp.htm.

[794] US Air Force. Safety investigations and reports. Technical report, US Air Force Departmental Publishing Office, Washington DC, USA, 1999. http://afpubs.hq.af.mil.

[795] US Army. Safety Policies and Procedures for Firing Ammunition for Training, Target Practice and Combat. Technical Report Army Regulation 38563 MCO P3570.1A, Headquarters Department of the Army and Navy, Washington, DC, 1983. http://www.usapa.army.mil/pdffiles/r385_63.pdf.

[796] US Army. Army Accident Investigation and Reporting. Technical Report Pamphlet (PAM 385-40) Safety, Headquarters Department of the Army, Washington, DC, 1994. http://www.usapa.army.mil/pdffiles/p385_40.pdf.

[797] US Army. Decorations, Awards, and Honors Army Accident Prevention Awards Program. Technical Report Army Regulation 67274, Headquarters Department of the Army, Washington, DC, 1995. ftp://pubs.army.mil/pub/epubs/pdf/r672_74.pdf.

[798] US Army. Risk Management. Technical Report Field Manual FM 100-14, Headquarters Department of the Army, Washington, DC, 1998. http://www.adtdl.army.mil/cgi-bin/atdl.dll/fm/100-14.

[799] US Army Center for Army Lessons Learned. Coalition operations lessons learned database. Technical report, US Army, Fort Leavenworth, KS, USA, 1999. http://www.abca.hqda.pentagon.mil/COLL%20Database/armies.html.

[800] US Army Joint Readiness Training Center. Trends, Tactics, Techniques and Procedures: 4th Quarter Financial Year '99 & 1st Quarter 2000. Technical Report Intelligence BOS (TA.5) 01-6, US Army, Fort Leavenworth, KS, USA, 2000. http://call.army.mil/products/ctc_bull/01-6/ta5bos.htm.

[801] US Army, Office of the Assistant Secretary of the Army. A-76 Contracting Lessons Learned. Technical report, Department of the Army, Falls Church, VA, USA, 2000. http://acqnet.sarda.army.mil/acqinfo/lsnlrn/Ini_main.htm.

[802] US Army, Product Manager, Paladin/FAASV. Drivers' Hatch. *Paladin/FAASV Newsletters*, page 8, 2000. http://w4.pica.army.mil/paladin/news/2qfy00.pdf.

[803] US Army Safety Centre. Accident investigation handbook. Technical report, US Army, Fort Rucker, Alabama, USA, 1999. http://safety.army.mil/pages/investigation/.

[804] US Army Safety Centre. Accidents kill more than the enemy. *Countermeasure*, 20(10), 1999. Available from http://safety.army.mil.

[805] US Army Safety Centre. Commander and staff risk management booklet. Technical report, US Army, Fort Rucker, Alabama, USA, 1999. Available from http://safety.army.mil/.

[806] US Army Safety Centre. Accident investigation and reporting procedures handbook. Technical report, US Army, Fort Rucker, Alabama, USA, 2000. http://safety.army.mil/pages/data/hazards/accidents/accident.pdf.

[807] US Army Safety Centre. Accident investigation: The other side of risk management. *Countermeasure*, 22(1), 2001. Available from http://safety.army.mil.

[808] US Army Safety Centre. Correction M939A2 Trucks. *Countermeasure*, 22(7), 2001. Available from http://safety.army.mil.

[809] US Army Safety Centre. An M551A1 in the wrong hands. *Countermeasure*, 22(2), 2001. Available from http://safety.army.mil.

[810] US Army Safety Centre. New Requirements for the M9 ACE. *Countermeasure*, 22(6), 2001. Available from http://safety.army.mil.

[811] US Army Safety Centre. No Recall On IPFUs. *Countermeasure*, 22(6), 2001. Available from http://safety.army.mil.

[812] US Army Safety Centre. The rest of the story. *Countermeasure*, 22(5), 2001. Available from http://safety.army.mil.

[813] US Army Safety Centre. Safety notice: Confidence course obstacle. Technical report, US Army, Fort Rucker, Alabama, USA, 2001. http://safety.army.mil, Note SAN 210800MAR00.

[814] US Army Safety Centre. Safety notice: HMMWV seatbelt. Technical report, US Army, Fort Rucker, Alabama, USA, 2001. http://safety.army.mil, Note SAN 200010181500.

[815] US Army Safety Centre. Safety notice: Tents and heaters. Technical report, US Army, Fort Rucker, Alabama, USA, 2001. http://safety.army.mil, Note SAN071445ZMAR0.

[816] US Army Safety Centre. A turn for the worse, M1A1. *Countermeasure*, 22(1), 2001. Available from http://safety.army.mil.

[817] U.S. Army Technical Center for Explosives Safety: D. Ford. Accident reporting: Do it right! *Explosives Safety Bulletin*, 11(3), 2000. http://www.dac.army.mil/ES/esm/bulletin/Vol11,iss3.PDF.

[818] U.S. Army Technical Center for Explosives Safety: D. Ford. Last minute change is fatal: how could this happen? *Explosives Safety Bulletin*, 12(1), 2000. http://www.dac.army.mil/es/esm/bulletin/Vol12,iss1.PDF.

[819] U.S. Army Technical Center for Explosives Safety: G. Heles. Recent HQDA Explosives Safety Policy Actions. *Explosives Safety Bulletin*, 11(3), 2000. http://www.dac.army.mil/ES/esm/bulletin/Vol11,iss3.PDF.

[820] U.S. Army Technical Center for Explosives Safety: R. Durand and D. Ford. Cook-offs. *Explosives Safety Bulletin*, 11(3), 2000. http://www.dac.army.mil/ES/esm/bulletin/Vol11,iss3.PDF.

[821] US Aviation and Missile Command, Safety Office. Preliminary incident notification form. Technical report, US Army, Redstone Arsenal, Alabama, USA, 2000. http://www.redstone.army.mil/safety/report.html.

[822] US Bureau of the Census. Detailed language spoken at home and ability to speak English for persons 5 years and over. Technical Report CPHL-133, Education and Social Stratification Branch, Population Division, US Bureau of the Census, Washington DC, United States of America, 1990. http://www.census.gov/population/socdemo/language/table5.txt.

[823] US Bureau of the Census. Home computers and Internet use in the United States: August 2000 special studies. Technical Report P23-207, U.S.Department of Commerce Economics and Statistics Administration, US Bureau of the Census, Washington DC, United States of America, 2000. http://www.census.gov/prod/2001pubs/p23-207.pdf.

[824] US Central Command. Investigation into the live-fire incident involving a U.S. Navy F/A-18 aircraft that dropped three 500-pound bombs on Observation Post 10 at the Udairi Range, Kuwait, 12 March 2001. Technical report, USCENTCOM, MacDill Air Force Base, Fla, USA, 2001. http://www.centcom.mil/kuwait/report/ud_range_completed.pdf.

[825] US Coast Guard. United States Coast Guard vs. Merchant Mariner's Document NO. Z-529214 License NO. 453652. Technical Report 2080, USCG, Washington DC, United States of America, 1976. http://www.uscg.mil/hq/g%2Dcj/appeals

[826] US Coast Guard. Investigation into the Circumstances Surrounding the Collision Between the Passenger Vessel Noordam and the Freight Vessel Mount Ymitos. Technical Report 16732, USCG, Washington DC, United States of America, 1995. http://www.uscg.mil/hq/g%2Dm/moa/omao1a5.htm.

[827] US Coast Guard. Portable Space Heaters. Technical Report Safety Advisory 0396, USCG, Washington DC, United States of America, 1996. http://www.uscg.mil/hq/g-m/moa/docs/sa0396.htm.

[828] US Coast Guard. Report of the Coast Guard-American waterways Organisation Quality Action Team on Towing Vessel Crew Fatalities, July 26, 1996. Technical report, USCG, Washington DC, United States of America, 1996. http://www.uscg.mil/hq/g-m/moa/docs/cafata.htm.

[829] US Coast Guard. Sea Kayak Safety Advisory. Technical Report Safety Advisory 0796, USCG, Washington DC, United States of America, 1998. http://www.uscg.mil/hq/g-m/moa/docs/sa0796.htm.

[830] US Coast Guard. International Maritime Information Safety System. Technical report, USCG, Washington DC, United States of America, 2000. http://www.uscg.mil/hq/g-m/moa/xnearm.htm.

[831] US Coast Guard. Investigation report into the circumstances surrounding the loss of the commercial fishing vessel ADRIATIC eight nautical miles East of Barnegat Light, New Jersey on January 18, 1999, with the loss of four lives. Technical Report 16732/MC99000698, USCG, Washington DC, United States of America, 2000. http://www.uscg.mil/hq/g-m/moa/docs/adriatic.pdf.

[832] US Coast Guard. Office of Investigations and Analysis: Safety Alerts and Lessons Learned. Technical report, USCG, Washington DC, United States of America, 2001. http://www.uscg.mil/hq/g-m/moa/safea.htm.

[833] US Coast Guard. Report into the loss of the commercial fishing vessel CAPE FEAR three nautical miles Southwest of Cuttyhunk, Massachusetts on January 9, 1999 with the loss of two lives. Technical report, USCG, Washington DC, United States of America, 2001. http://www.uscg.mil/hq/g-m/moa/docs/capefear.pdf.

[834] US Coast Guard (S. Ferguson, J. Gelland, G. Kraatz, A. Landsburg, R. Nikas, C. Pillsbury, A. Rothblum). International Maritime Information Safety System - Project Blueprint. Technical report, USCG, Washington DC, United States of America, 2000. http://www.uscg.mil/hq/g-m/moa/docs/blue.htm.

[835] US National uesearch Council, Subcommittee on Coordinated Research and Development Strategies for Human Performance to Improve Marine Operations and Safety. *Advancing The Principles of The Prevention Through People Program.* National Accademy Press, Washington, DC, USA, 1997. http://books.nap.edu/books/NI999999/html/28.html#pagetop.

[836] US Navy, Chief of Naval Operations. Third Endorsement of the Investigation to Inquire into the Actions of USS Cole in Preparing for and Undertaking a Brief Stop for Fuel at Bandar at Tawahi (Aden Harbour) Aden, Yemen, on or About 12th October 2000. Technical report, Department of the Navy, Washington, DC, USA, 2000. http://www.foia.navy.mil/USSCOLE/ENDORSE/CNO.pdf.

[837] US Navy, Commander in Chief, US Atlantic Fleet. Second Endorsement of the Investigation to Inquire into the Actions of USS Cole in Preparing for and Undertaking a Brief Stop for Fuel at Bandar at Tawahi (Aden Harbour) Aden, Yemen, on or About 12th October 2000. Technical report, Department of the Navy, Washington, DC, USA, 2000. http://www.foia.navy.mil/USSCOLE/ENDORSE/CINCLANTFLT.pdf.

[838] U.S. Nuclear Regulatory Commission. Reactor safety study. Technical report, NUREG, Washington, DC, USA, 1974.

[839] R. Usher and V. Jakupec. *Policy Matters: Flexible Learning and Organisational Change.* Routledge Studies in Human Resource Development. Taylor and Francis, London, United Kingdom, 2001.

[840] T.W. van der Schaaf. *Near Miss Reporting in the Chemical Process Industry.* PhD thesis, Technical University of Eindhoven, Eindhoven, The Netherlands, 1992.

[841] T.W. van der Schaaf. PRISMA: A risk management tool based on incident analysis. In *International Workshop on Process Safety Management and Inherently Safer Processes*, pages 242–251, Orlando, Florida, USA, October 8-11 1996.

[842] T.W. van der Schaaf. Prevention and recovery of errors in software systems. In C.W. Johnson, editor, *Proceedings of the 1st Workshop on Human Error and Systems Development*, GAAG TR-97-2, Glasgow, Scotland, 1997. Glasgow Accident Analysis Group, University of Glasgow. http://www.dcs.gla.ac.uk/ johnson/papers/workshop/human_error.htm.

[843] T.W. van der Schaaf, D.A. Lucas, and A.R. Hale. *Near Miss Reporting as a Safety Tool.* Butterworth-Heinemann, Oxford, United Kingdom, 1991.

[844] W. van Vuuren. *Organisational Failure: An Exploratory Study in the Steel Industry and the Medical Domain.* PhD thesis, Institute for Business Engineering and Technology Application, Technical University of Eindhoven, Eindhoven, The Netherlands, 2000.

[845] T. Vara. Female soldier readiness: A leader's guide. Technical report, US Army Center for Army Lessons Learned and Madigan Army Medical Center, Tacoma, WA, USA, 2000. Available from http://call.army.mil.

[846] D. Vaughan. *The Challenger Launch Decision.* Chicago University Press, Chicago, United States of America, 1996.

[847] S.A. Viller. *Human Factors in Requirements Engineering: A Method for Improving Requirements Processes for the Development of Dependable Systems.* PhD thesis, Department of Computer Science, University of Lancaster, Lancaster, United Kingdom, 2000.

[848] C. Vincent, S. Taylor-Adams, J. Chapman, D. Hewett, S. Prior, P. Strange, and A. Tizzard. How to investigate and analyse clinical incidents: Clinical risk unit and association of litigation and risk management protocol. *British Medical Journal*, 320(7237):777–781, 2000.

[849] C. Vincent, S. Taylor-Adams, and N. Stanhope. Framework for analysing risk and safety in clinical medicine. *British Medical Journal*, pages 1154–1157, 1998.

[850] J. W. Vincoli. *Basic Guide to Accident Investigation and Loss Control.* Van Nostrand Reinhold, New York, USA, 1994.

[851] Volpe Research Centre. Safety. Technical report, US Department of Transportation, Cambridge, MA, United States of America, 2001. http://www.volpe.dot.gov/infosrc/journal/30th/safety.html.

[852] W.A. Wagenaar and J. Groeneweg. Accidents at sea : Multiple causes and impossible consequences. In E. Hollnagel, G. Mancini, and D.D. Woods, editors, *Cognitive Engineering In Complex Dynamic Worlds*, pages 133 – 144. Academic Press, London, United Kingdom, 1988.

[853] W.A. Wagenaar, J. Groeneweg, P.T.W. Hudson, and J.T. Reason. Promoting safety in the oil industry. *Ergonomics*, 37(12):1999–2013, 1994.

[854] B. Wallace, A. Ross, and J. Davies. Proposed root-case codings based on a new event causal framework: Interim report on results of inter-rater and intra-rater reliability. Technical report, Department of Applied Psychology, University of Strathclyde, Glasgow, UK, 1999. Report for British Nuclear Fuels, Cited in [197].

[855] Washington Metro Area Transit Authority. Metro reports a decline in customer injuries in the first quarter of fiscal year 2002. Technical report, WMATA, Washington DC, USA, February 2002. http://www.wmata.com/about/MET_NEWS/200202/pr_safety_performance.htm.

[856] P. Wasielewski. Speed as a measure of driver risk: Observed speed versus driver and vehicle characteristics. *Accident Analysis and Prevention*, pages 89–103, 1984.

[857] R. Webb. Issues in rail reform. Technical Report Research Paper 14 1999-2000, Economics, Commerce and Industrial Relations Group, Parliament of Australia, Canberra, Australia, 2000. http://www.aph.gov.au/library/pubs/rp/1999-2000/2000rp14.htm.

[858] M. Wellbank. An overview of knowledge acquisition methods. *Interacting with Computers*, 2(1):83–91, 1990.

[859] G.L. Wells. Eye witness testimony. In A. E. Kazdin, editor, *Encyclopedia of Psychology.* American Psychological Association/Oxford University Press, Oxford, UK, 1999.

[860] G.L. Wells, R.C. lindsay, and T.I. Ferguson. Accuracy, confidence and juror perception in eyewitness testimony. *Journal of Applied Psychology*, pages 440–448, 1979.

[861] G.L. Wells, M. Small, S.J. Penrod, R.S. Malpass, S.M. Fulero, and C.A.E. Brimacombe. Eyewitness identification procedures: Recommendations for lineups and photospreads. *Law and Human Behaviour*, pages 603–647, 1998.

[862] R. Westrum. Cultures with requisite imagination. In J. Wise, D. Hopkin, and P. Stager, editors, *Verification and Validation of Complex Systems: Human Factors Issues*, pages 401–416. Springer Verlag, Berlin, Germany, 1992.

[863] C.D. Wickens. *Engineering Psychology and Human Performance.* Harper Collins, New York, NY, United States of America, 1992. Second Edition.

[864] E. L. Wiener. Cockpit automation. In E. L. Wiener and D.C. Nagel, editors, *Human Factors in Aviation*, pages 433–461. Academic Press, San Diego, CA, United States of America, 1988.

[865] G.J.S. Wilde. The theory of risk homeostasis: Implications for safety and health. *Risk Analysis*, 2:209–225, 1988.

[866] J.A. Williamson, R.K. Webb, A. Sellen, W.B. Runciman, and J.H. Van der Walt. The Australian Incident Monitoring Study. human failure: an analysis of 2000 incident reports. *Anaesthesiology Intensive Care*, 21(5):678–683, 1993.

[867] M. A. Wollerton. Last printing of user facility reporting bulletin. *FDA User Facility Reporting Bulletin*, 1997.

[868] W.B.L. Wong, P.J. Sallis, and D. O'Hare. Information portrayal requirements: Experiences with the critical decision methodology. In H. Thimbleby, B. O'Conaill, and P. Thomas, editors, *People and Computers XII: Proceedings of HCI'97*, pages 397–415. Springer Verlag, London, United Kingdom, 1997.

[869] R.H. Wood and R.W. Sweginnis. Analytical techniques for mid-air collision investigations. In P. d'Antonio and C. Ericson, editors, *Proccedings of the 16th Annual Systems Safety Conference*, pages 346–351, Unionville, VA, United States of America, 1998. Systems Safety Society.

[870] K. Woodcock and A. Smiley. Developing simulated investigations for occupational accident investigation studies. Technical report, School of Occupational and Public Health, Ryerson University, Toronto, Canada, 2001.

[871] D.D. Woods. Cognitive demands and activities in dynamic fault management. In N. Stanton, editor, *The Human factors of Alarm Design*, pages 63–92. Taylor and Francis, London, United Kingdom, 1994.

[872] World Health Organisation. *Health and Environment in Sustainable Development: Five Years after the Earth Summit.* WHO, Geneva, Switzerland, 1997. http://www.who.int/environmental_information/Information_resources/htmdocs/execsum.htm.

[873] World Health Organisation. Occupational health: Ethically correct, economically sound. *WHO Fact Sheets*, 84, 1999. http://www.who.int/inf-fs/en/fact084.html.

[874] L. Wright. Towards an empirical test of the iceberg model. In P.C. Cacciabue, editor, *Human Decision Making and Manual Control: EAM2000*, EUR 19599 EN, pages 145–152. European Commission, Joint Research Centre, Ispra, Italy, 2000.

[875] P. Wright and A. Lickorish. Proof reading texts on screen and paper. *Behaviour and Information Technology*, 2(3):227–235, 1983.

[876] P.C. Wright and A.F. Monk. A cost-effective evaluation method for use by designers. *International Journal of Man-Machine Studies*, 35(6):891–912, 1991.

[877] A.W. Wu. Medical error: the second victim. *British Medical Journal*, 320(7237):726–727, 2000.

[878] N. Yoshimura and P. Anderson. *Inside the Kaisha : Demystifying Japanese Business Behaviour*. Harvard Business School Press, Harvard, MA, USA, 1997.

[879] H. Zerkani and R. Dumolo. System safety lifecycle based on IEC61508 and its use for railway applications. In P. d'Antonio and C. Ericson, editors, *Proccedings of the 16th Annual Systems Safety Conference*, pages 369–379. Systems Safety Society, Unionville, VA, United States of America, 1998.

[880] M. Zviran and W. Haga. A comparison of password techniques for multilevel authentication systems. *Computer Journal*, 36(3):227–237, 1993.

Index