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 [774]. 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

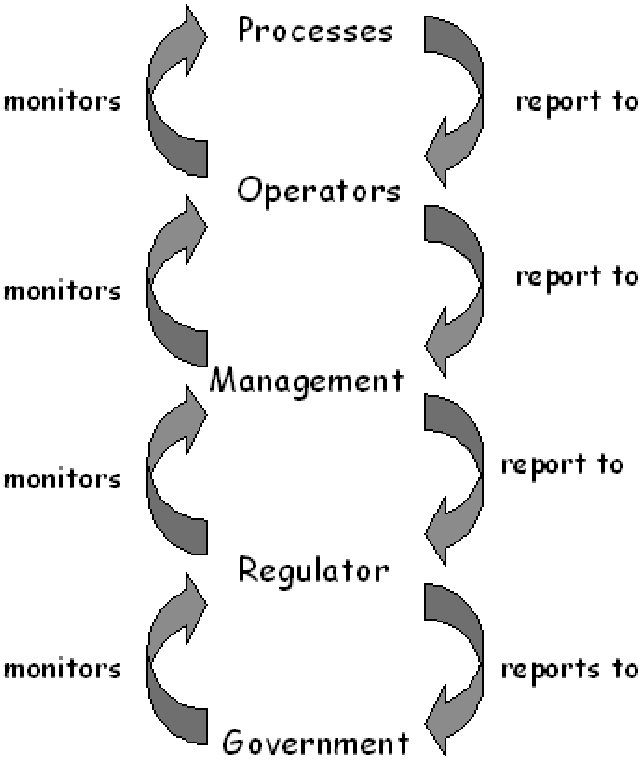

Figure 3.1: Levels of Reporting and Monitoring in Safety Critical Applications

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 [702]. 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 [691]. 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 [863]. 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].

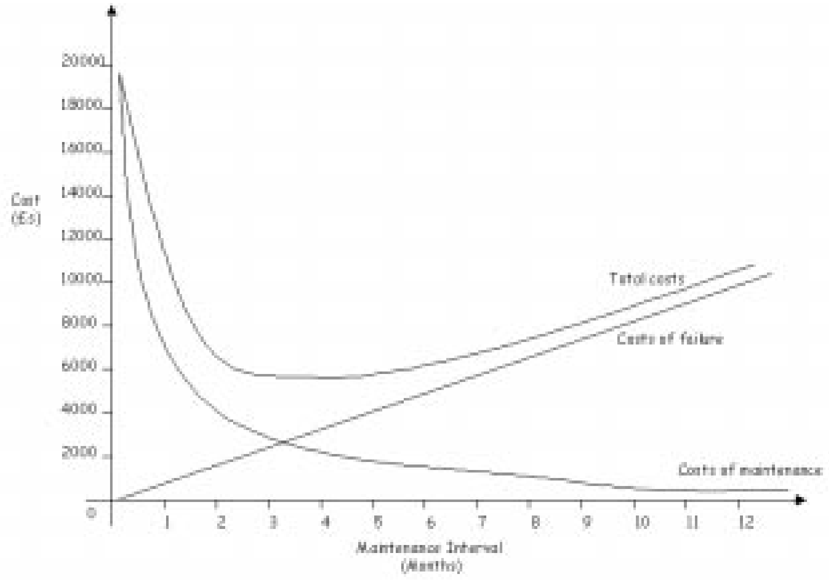

Figure 3.2: Costs Versus Maintenance Interval

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 [763]. 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

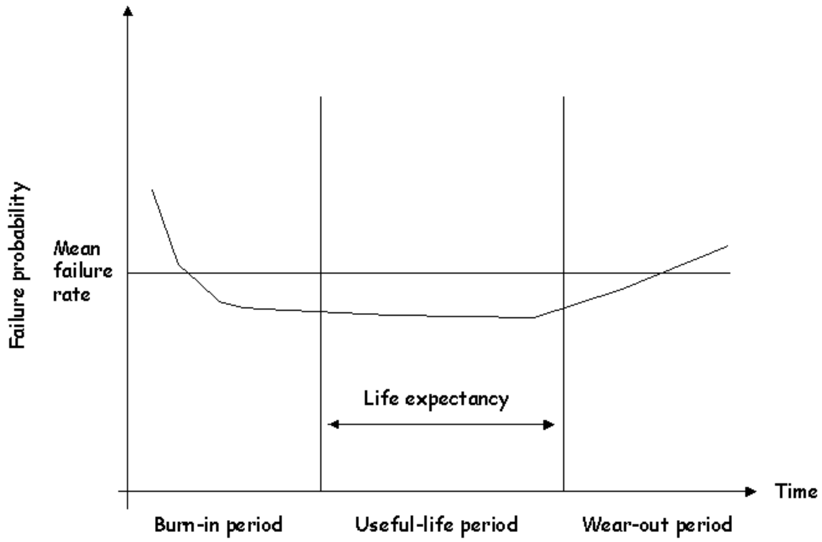

Figure 3.3: Failure Probability Distribution for Hardware Devices

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. ([773], 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 [628].

- *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." ([773], 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." [773]

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, [773]).

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 [880]. 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 [700] and Wickens [864] 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" [767] 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 [864].

- *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 [761]. Mental workload is more problematic. Wickens identifies a number of key questions about workload that can be adapted to guide incident investigation [864]. 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 [771]. 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 [713]. 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 [723]. Higher speeds have been observed for drivers who have a previous record of accidents [857].

> Rockwell's pioneering study showed that electrical workers who take higher risks in their daily lives are also involved in more accidents at work [712]. 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 [772]. 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 [864]. 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 [864]. 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 [652]. 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]. However, many physiological problems also stem from the behaviour of the workers themselves. Workers in many industries, including car production, marine engineering and electricity generation, have been shown not to engage in risk reducing behaviour [311]. Instead, 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 dangers of physiological injury [77]. Incident reporting systems, such as that proposed by the OSHA standard, have been advocated as a means of addressing these problems. Specific instances of injuries through inappropriate posture can be used to reinforce the potential consequences of violations. Direct information about real incidents often proves to be more effective than abstract classroom-based training sessions [425].

Perceptual failures describe incidents in which human operators fail to correctly detect important cues and signals in our environment. As before, it is possible to identify a number of causal factors that are common to many of these occurences. Poor design also plays a role in creating incidents that are related to perceptual failure. Many authors have commented on the clutter that characterises many cockpits [864]. The crews' apparent inability to sample their information sources has been identified as a causal factor in numerous accidents, most notably the Kegworth crash [8]. However, as Billings notes, there is often a tension between filtering information to reduce the perceptual loading on operators and actively hiding information that may be essential for fault diagnosis and other control tasks [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 [739]. At the other end of the scale are systems that provide their operators with almost no perceptual cues about a potential failure. 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]

The experienced physicians who had set up this device had assembled it so that it allowed a free flow of the drug into the patient once the physical barrier of the stopcock was removed. The device was started but there was no flow of the drug because the stopcock was closed and so a visual and an auditory alarm were presented. When the stopcock was opened, the device again failed to detect any drops of the drug being administered and the same alarms were presented. In this case, the device could not detect drops being administered because the drug was passing freely into the patient. The 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.

Environmental factors also 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 [669]. 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." [668]

The previous paragraphs describe how poor design and environmental features, such as background noise, can impair an operator's ability to perceive critical information. These perceptual problems can be analysed in more detail. Many failures relate narrowly to the problems of signal detection. Table 3.5.1 explains some of the issues involved in this aspect of perception. As can be seen, 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 3.5.1. 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. From this it follows that many of the perceptual problems in incident reports are either missed signals or false alarms. It might, at first sight, appear that a false alarm should not jeopardise the safety of the system. However, things are often less clear cut than table 3.5.1 suggests. There have been situations in which the response to such a 'non-event' has trigered a real incident [864]. 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].

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

Table 3.2: Outcomes from Signal Detection

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 conducted numerous studies, in domains including steel production and healthcare, that identify the different problems that arise from both rapid and gradual changes in the presentation of information [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 [686]. Both audio and proprioceptive alarms provoke faster responses than visual warnings. However, visual information rapidly captures the operator's attention. Response times are slower for visual stimuli but they given an audible and a visual warning operators will more reliably provide the response associated with the visual rather than the audio alarm. 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.

Much more could be said about the ways in which the human perceptual system contributes to, and helps to avoid, major incidents. Wickens provides an excellent overview of this area [864]. However, it is worth emphasising that perception cannot be isolated from other attributes of human behaviour. In particular, our sampling behaviour is heavily determined by cognitive or mental models of the processes being observed. People will sample channels with higher error rates more frequently that those with lower error rates. Unfortunately, our internal stochastic models of our environment are not updated as often as they might be. As a result, we do not adjust our sampling rates to reflect changes in application processes. 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 [738]. 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 that 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. The previous paragraph has also argued that 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.
>
> The Safety Board determined from the stopped position of the MARC train 286 locomotive and its event recorder information that the engineer placed the train into emergency braking 1,407 feet before the collision at a speed of about 66 mph. The engineer made the emergency brake application about 510 feet after passing the optimum sight distance location, about 1,227 feet from the EAS-2 or 5.27 seconds later. The delay is understandable and reasonable considering the engineers apparent belief that he was

> operating under a CLEAR signal indication.
>
> There is no reason to suppose that the MARC train 286 engineer would be looking for the Georgetown Junction signal as soon as it was physically visible. 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." [597]

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 [772]. 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,

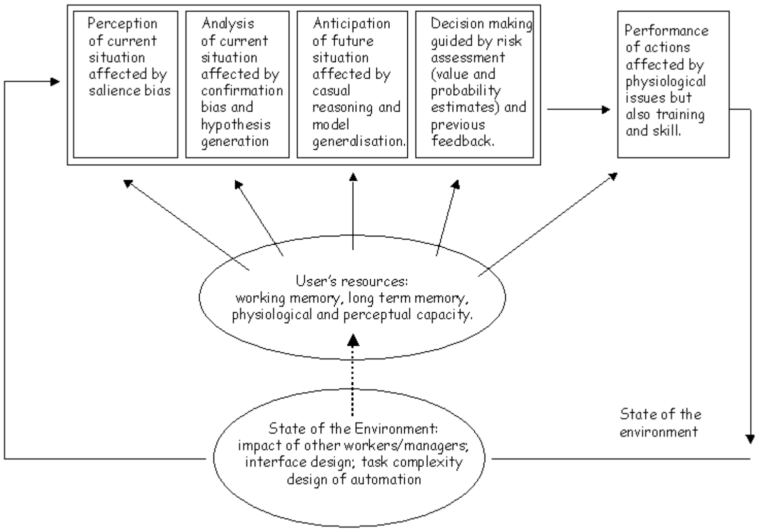
Perception of current situation affected by salience bias
Analysis of current situation affected by confirmation bias and hypothesis generation
Anticipation of future situation affected by casual reasoning and model generalisation.
Decision making guided by risk assessment (value and probability estimates) 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 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:

> "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].

Confirmation bias is, in turn, directed by our ability to consider a number of competing hypotheses. For example, there is evidence that many people 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 [864]. 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]. This work confirms previous empirical evidence that has great significance for the effective operation of incident reporting systems. 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 the perception and analysis of the current situation are also closely tied to our 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. This, in turn, can lead to incidents. 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 a particular inhalational device led to incorrect predictions about the induction of an inhalational 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%." [756]

Norman [637] again illustrates the connections between cognition and perception when he argues that the development of appropriate models can be supported by the provision of appropriate feedback about system behaviour. In this case, more prominent information about the particular device characteristics would have supported the users' already adequate mental model about the general operation of devices of this sort.

Figure 3.4 illustrates further ways in which cognition influences the actions of human operators. In particular, it illustrates how decision making is again closely 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]

It is important to note that figure 3.4 does not show a linear progression between the perception of the current situation, the analysis of that situation, the prediction of future situations and eventual decision making. This is a significant weakness of many previous models. Operators will iterate between these phases before taking actions. Conversely, there are incidents in which it seems that some of these phases are completely ignored.

The previous figure 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. However, other factors such as managerial or domestic pressures may act to influence operator behaviour in a less direct manner.

Previous sections have argued that attributes of human cognition, perception and physiology play an important role in many incidents and accidents. The following section builds on this analysis by focussing on the specific ways in which human error and violations can jeopardise the safety of many complex systems.

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]. Again these different control mode 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. It is important also to emphasise that violations may actually 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 employed by the North Anna reactor led to dangerous temperature profiles following a scram. The operators were faced with a difficult choice. Following the Three Mile Island accident, 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].

It is possible to distinguish between unintended and deliberate violations. If an individual does not know that they are violating a rule or procedure then this can be interpreted as an error. Unfortunately, the pragmatic consequence of such theoretical distinctions is that an incident investigator must be able to accurately discern the intentions of an operator. For now it is sufficient to focus on deliberate, or knowing, violations. Later chapters will return to the problems of identifying intentions from reports about adverse events. 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 [702]. Reports of such incidents are particularly instructive because they illustrate situations in which rules and regulations may actually place staff in danger. The usual emphasis of compliant action might, under other circumstances, have led to a much worse outcome than the one that was reported. 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 possible 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' [702].

Unfortunately, experience suggests that many incidents occur because of more complex combinations of optimising, necessary and routine violations. This is illustrated by a report that was issues by the US Chemical Safety and Hazard Identification Board into an incident at an explosive company:

> "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.

This analysis also reveals that 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.

The previous example illustrates the general problems that arise when analysing the nature of violations that contribute to incidents and accidents. As noted in the previous chapter, 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. However, it is also important to note that in the case of optimising violations, it is likely that some sections of management may actually collude in the breaking of rules and regulations. In such circumstances, the reporter (whistle blower) must not only be assured of their anonymity but also of the independence of any subsequent investigation.

Just as it is possible to 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 and lapses 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. A slip often has visible consequences, such as a slip of the tongue, in which it is possible to observe that an error has occurred. A lapse describes more covert forms of error, including failures of memory such as forgetting someones name, that may only be apparent to the person experiencing them. Both of these error forms can be distinguished from mistakes which, as we have seen, relate to failures of intention rather than execution. Mistakes stem from a failure to select appropriate objectives irrespective of whether or not the actions taken to achieve those objectives are successful.

Reasons Generic Error Modelling (GEMS) approach is one of a number of extensions to the slip, lapse and mistake taxonomy [700] GEMS is heavily influenced by Rasmussen's Skill, Knowledge and Rules approach to cognition [695]. These represent different levels of performance. 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. Rule based and knowledge 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, operator may incorrectly identify the problem at hand and, therefore, select rules and procedures that are more appropriate to another problem. Alternatively, users may apply the wrong rules and procedures that are applicable to a situation which they have correctly diagnosed. In other words, 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 are also typical of incidents that involve encysting; operators will continue to focus in munite detail at some small aspect 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. It, therefore, follows that Reason also focuses on the monitoring failures that may prevent an operator from effectively instigating 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. It is the failure of these attentional checks that are liable to 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 system feedback.

There are a number of justifications for introducing the distinctions between rule, skill and knowledge based performance in addition to the distinctions between slips, lapses and mistakes. These terms also playing an increasingly important role in the techniques that are used to analyse safety-critical incidents. For instance, slips, lapses and mistakes are all included within EUROCONTROL’s harmonisation of European Incident Definitions Initiative for Air Traffic Managment [718]. This is developing 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)

Although the statistics cited for the human contribution in incidents are impressive, it is perhaps even more surprising that human error does not play an even larger role than it already does. People continually make mistakes, commit slips or suffer from lapses of attention. Very few of the errors and violations that we commit will ever result in an incident or accident. Even less result in an official or confidential report. This apparent paradox is explained by the monitoring activities that were mentioned in previous paragraphs. We regulate our behaviour to ensure that we minimise our chances of paying the potential costs of erroneous acts and violations. 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 upon the genotypes and phenotypes of individual human error. Little attention has been paid to the particular problems of coordinating interaction with other members of a working group or team. In contast, Viller [848] provides a summary of social and group performance failures:

- failures due to distraction . For example, where an individual interrupts one of their colleague’s tasks.

- failures due to performance effects. For example, individuals may consistently perform below expectations if they are worried about their actions being monitored or observed by their colleagues.
- failures that are due to inappropriate human resources in the group. For example, there is no competent group member.
- socio-motivational failures. For example, there may be 'free-riders' or others who mask individual poor performance by over-relying upon the performance of their colleagues in the group.
- group coordination failures. For example, the overhead of coordinating group actions can impair the effectiveness of the group as a whole.
- status related failures. For example, where the status of a group member mitigates against their contributions being taken as 'seriously' as they merit. Alternatuvely, a group may grant undue attention to other individuals.
- group planning and management failures. For example, groups may create unnecessary subtasks or may allocate them to inappropriate individuals.
- failures due to inappropriate leadership style. This analysis is based on the idea that there are two important styles of leadership. 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. For example, the appointed leader may not have the necessary skills that contribute to both of the roles mentioned above.
- failures due to excessive influence of the leader. For example, a high status leader may stifle contrary opinions in situations where they are, themselves, in the wrong.
- failures due to conformity arising from inappropriate normative influence. For example, when an incorrect judgement from a high status member commands influence because other respect that status rather than the value of the judgement itself.
- failures due to conformity arising from inappropriate informational influence. For example, when the judgement of one member is based on false evidence or is misunderstood by another member of the group.
- failures due to group polarisation and groupthink. For example, 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 quotation provides a concrete illustration of the incidents that stem from communications failures between the operators of complex 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 AFR 813. However, Departures did not inform the Arrivals officer of another aircraft BAW 818 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 AFR 813 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 SAB 603 to the right to achieve maximum separation with AFR 813 and also to minimise any disruption to the latter aircraft's flightpath. This resulted in SAB 603 and BAW 818 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 BAW 818 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]

Such incidents are instructive because they typify 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. In the following discussion, it is important not to forget that 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 [486].

It is important not to underestimate the problems that arise when attempting to understand the deeper causes of team-based failures [729]. At the most superficial level, it is possible to view these genotypes as simple elaborations of the single-person failures that were examined in previous pages. 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. However, there are many ways in which Figure 3.5 represents a gross simplification. For example, there are also

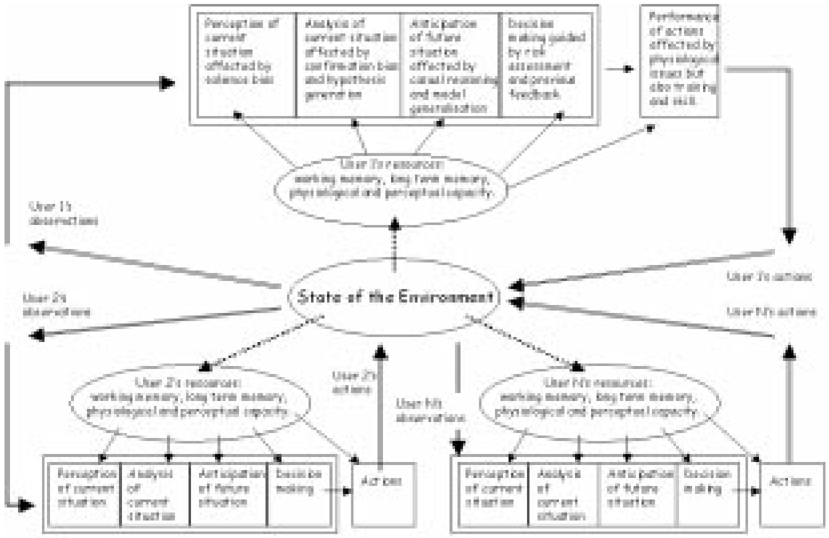
Perception of current situation affected by salience bias
Analysis of current situation affected by confirmation bias and hypothesis generation
Anticipation of future situation affected by causal reasoning and model generalisation
Decision making guided by risk assessment and previous feedback
Performance 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
State of the Environment
User 1's actions
User 2's observations
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

ways in which 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 groups that exhibit a more 'precautionary principle'. These teams 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. However, if individuals initially accept higher risk positions then the group is liable to adopt even higher risk decisions.

There are further aspects of group interaction that are not directly captured by Figure 3.5. In particular, a number of studies have pointed to the incidents that can occur when teams make inefficient use of the personnel that are available to them. For instance, the following quotations are taken from an investigation by the Transportation Safety Board of Canada into incidents of 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 TSB 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 OOW's (Officer of the Watch) 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 OOW was not monitoring the pilot's actions and did not recognize that the change of course was premature. The OOW appeared to have placed total confidence in the pilot's navigation ability. When the pilot passed his position report to VTS, the OOW logged the time, but he did not plot the position on the chart. Had the OOW 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)" [620]

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 [848]. 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 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:

> "Several of the firefighters who responded to the accident stated that they received contradictory information from Metrorail personnel at the scene when they asked if the third rail had been deenergized. According to recordings of tower communications,

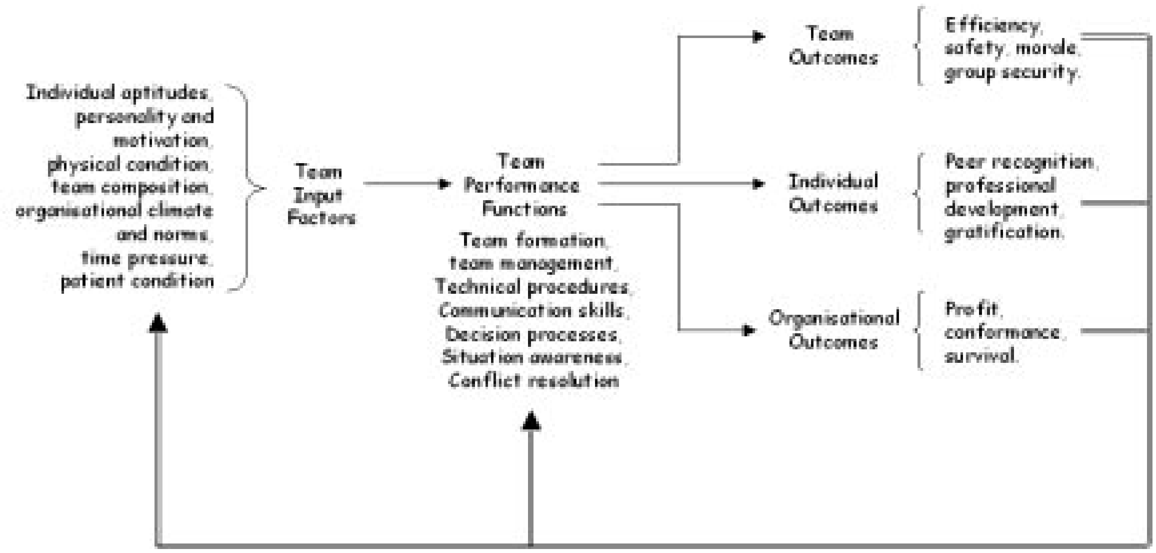
Individual aptitudes, personality and motivation, physical condition, team composition, organisational climate and norms, time pressure, patient condition
Team Input Factors
Team Performance Functions
Team formation, team management, Technical procedures, Communication skills, Decision processes, Situation awareness, Conflict resolution
Team Outcomes
Efficiency, safety, morale, group security.
Individual Outcomes
Peer recognition, professional development, gratification.
Organisational Outcomes
Profit, conformance, survival.

Figure 3.6: Influences on Group Performance

> Metrorail personnel at the scene contacted the yard tower on several occasions to request that power at the accident scene be brought down... After yard third-rail power had been brought down, the tower operator replied to requests from the scene by informing callers that yard power was down but that the Operations Control Centre controlled power in the area of the accident. On at least one occasion he warned the caller that rescuers should hotstick the gaps before proceeding with any work [ie, test the 3rd rail with a volt probe device to see if it is energized] There was no evidence that a direct communication link was ever established between firefighters on the scene and OCC personnel.” [591]

.

3.6.1 Common Ground and Group Communication

Grice [296] has developed a number of guidelines that are intended to support communication with groups of co-workers:

1. Be as informative as is required but not more so

2. Say what is true, not that for which you lack sufficient evidence

3. Be relevant

4. 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 thier 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. This can best be illustrated by part of a transcript from a cockpit voice recorder imediately before the crew shut-down their one healthy engine:

> “From the CVR 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 ATC 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 ATC 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]

As usual with complex incidents of human error, it is possible to identify multiple hypotheses about the causes of this error. It might be that the events between the crew's initial conversation and the First Officer's action interfered with the First Officer's recollection of what had been decided. Alternatively, however, one can look more closely at the transcript to identify a communications problems between the crew. 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.

This simple analysis of common ground is not, however, a sufficient explanation of communication failure. In particular, it is important to understand why team members may fail to perform the cross-checking that may be necessary to ensure that they accurately understand the meanings behind their colleagues' utterances. One explanation for this is that establishing common ground will carry a number of potential costs. Table 3.6.1 lists some of overheads involved in refining our understanding of a converstion. This analysis is particularly important because it considers the way in which the costs of repairing a potential mistake may actually be perceived to be more costly than executing an action based on partical knowledge [864]. In other situations, very similar events can lead to entirely different team behaviours. For example, individuals may initiate ask further questions to

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

clarify their understanding of their colleagues' beliefs and intentions if that indidividual has received appropriate training (see below) 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.

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] http://www.open.gov.uk/aaib/gobmm.htm

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 [91] 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."

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 [727, 662, 872]. 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." [592]."

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 [734]. 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)" [595]

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 [734] have argued that crew coordination practices are ingrained more deeply if they are treated as a key component of many routine tasks [735]. 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 [848]. 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 [663].

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 [674]. 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.