Chapter 5

Detection and Notification

The previous chapter presented a number of different ways in which incident reporting systems can be organised. These architectures ranged from small-scale local systems through intermediate gatekeeper systems through to more complex, devolved, national and international mechanisms. The following chapters build on this by examining a number of generic problems that must be addressed by all incident reporting systems. These issues are illustrated in Figure 5.1. As can be seen, the

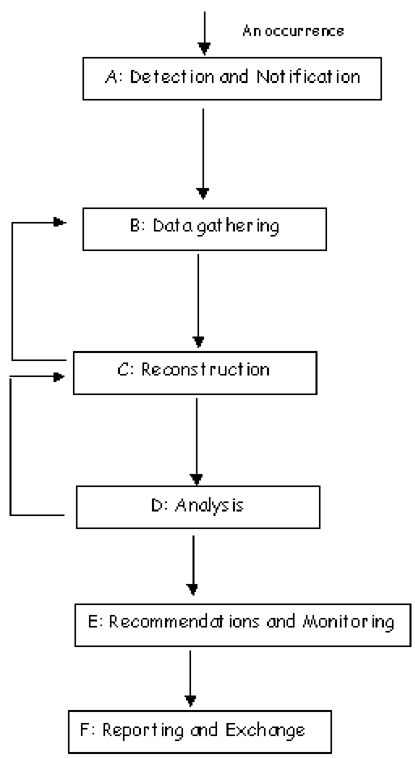

Figure 5.1: Generic Phases in Incident Reporting Systems

detection and notification of an occurrence is followed by a phase in which data is gathered about the events leading to a failure. This data can be used to reconstruct the likely ways in which events combined during the course of an incident. Once a probable reconstruction has been developed, it is possible to analyse these likely scenarios to identify key latent and catalytic causes. These form the focus for any subsequent recommendations about ways to prevent future failures. If these recommendations are adopted then they must be acted upon and their outcomes must be monitoried.

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

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

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

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

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

> "All staff, including contractors, should be regularly reminded of their responsibilities with regard to adverse incident reporting and of the relevant local procedures including the need to isolate and retain defective or suspect items. This information should also be conveyed to new staff as part of their induction training. The procedures should ensure that: where appropriate, a liaison officer is appointed with the necessary authority to take responsibility for the reporting of medical device related adverse incidents to the Medical Devices Agency (MDA) as detailed in the Annexes; devices involved in an adverse incident together with other material evidence (e.g., packaging of a single use device) should be clearly identified and kept in quarantine, where practicable, until MDA's device specialists have been consulted. Where quarantine is not practicable, the state of the device(s) at the time of the incident should be recorded for use in any subsequent investigation; local action is taken as necessary to ensure the safety of patients, users and o thers. Regular reviews should be undertaken to ensure that the procedures are effective and are being followed." [536]

As this quotation suggests, workers must receive training about what to report and how to report it. Setting up the necessary infrastructure for an incident reporting system does not guarantee that staff will be motivated to participate. This excerpt also stresses the importance of local liason officers, even in a large national reporting system. These trusted advocates support staff who are concerned about adverse occurrences. They must address contributors' concerns about anonymity and confidentiality that were described in Chapter 3.7 as part of a more general review of the key roles that support incident reporting systems. The net effect of these concerns is to exacerbate problems of under-reporting. Rather than reiterate the importance of addressing contributors' concerns about anonymity and confidentiality, the following paragraphs look at techniques that are specifically intended encourage the notification of adverse occurrences.

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

> "We know from previous studies that self-reporting of incidents retrieves only about 30% of the incidents reported by independent observer. We do not know, therefore, either true numerators nor because we do not collect them, denominators; even the Department of Health does not know how many anaesthetics are given annually. Any idea that this scheme might give absolute incident rates must therefore unfortunately be rejected. what we can hope to do is to paint a picture of what we are told nationally and allow departments to see whether the incidents that they are seeing locally are common or rare..." [716]

Vincent, Taylor-Adams and Stanhope observe that between 4-17% of patients in acute hospitals studies suffer from iatrogenic injury [850]. Observational studies have found that 45% of patients experienced some medical mismanagement and 17% suffered events that led to a longer hospital stay or more serious problems [28]. It has been estimated that approximately 850,000 adverse events occur within the UK National Health Service each year [635]. The earlier Harvard Medical Practice Study used similar techniques to estimate that among the 2,671,863 patients discharged from New York hospitals in 1984 there were 98,609 adverse events and 27,179 adverse events involving negligence [93]. Even the most successful voluntary reporting systems only succeed in eliciting information about a tiny fraction of the incidents that are revealed by the exhaustive analysis of records and logs. For instance, Barach and Small estimate that between 50 and 95% of medical incidents go unreported [66].

5.1.1 Reporting Bias

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

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

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

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

3. *automated incident detection* Clearly, the burdens of manually search for indications of incidents can frustrate attempts to obtain clear base-line measures of incident rates. As a result, a variety of automated tools (see below) can be used to search for key indicators. These tools range from simple databases through to more advanced data mining systems similar to those that will be discussed in Chapter 14.5. However, such tools introduce a further level of indirection that can bias results in ways that are often difficult to predict. In particular, there are the twin problems of precision and recall. A low precision search will detect many potential incidents that analysts must manually assess and then reject as not representing actual incidents. A low recall search will yield a number of potential incidents but will also leave many real cases undetected in the mass of incident data.

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

Jha, Kuperman, Teich, Leape, Shea, Rittenberg, Burdick, Segerand, Vander Vliet and Bates hae conducted several studies into the use of both manual and automated techniques for assessing base-line incident frequencies [402]. Most of their work focuses on adverse drug events which they argue are both common and costly. They criticise the 'spontaneous', voluntary systems in most hospitals as lacking sensitivity. They also criticise the costs associated with the exhaustive manual analysis of patient charts. As a result, they have worked to develop a computer-based adverse drug event monitor. Subsequent studies have then compared the performance of this tool with the products of both chart review and voluntary report systems. In one study, they focused on all patients admitted to nine medical and surgical units in an eight-month period [402]. The monitoring program identified situations that suggested a potential adverse drug event. These included requests for antidotes, such as naloxone. A trained reviewer then examined the patient's records to determine whether an adverse incident had occurred. The results were then compared with the products of an intensive manual review and a voluntary reporting system operated by nurses and pharmacists. Both the automated system and the chart review strategies were independent, and the reviewers were blinded.

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

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

> efficient approach for measuring ADE frequency and gauging the effectiveness of ADE prevention programs." [402]

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

> "In February 1999 a urologist at the Sturdy Memorial Hospital in Attleboro, Massachusetts, requested a retrospective review of a 1996 biopsy result because of the patient's clinical course and the results of a biopsy in 1999. The review revealed that the 1996 report was incorrect. The urologist and pathologist (neither of whom was responsible for the 1996 reading) implemented appropriate management for the affected patient.
>
> When they discovered a second misread prostate biopsy from the same period the urologist and pathologist became concerned that the frequency of these errors was higher than "expected". Fears about malpractice suits and damaged reputations emerged... Ultimately, the medical director thought that all the prostate biopsies performed during 1995-7, the period of tenure of the clinicians associated with the two errors, should be reviewed... During the review we wondered about any requirements to report to regulatory agencies. Our lawyers told us we had no obligation to report this kind of error... We decided to report our initial findings to the Department of Public Health and the Board of Registration in Medicine. In total 20 of the 279 prostate biopsies from 1995-7 were in error. The urologists caring for these 20 patients were told of the changes in the biopsy interpretations, and it was agreed that the urologists would contact each patient and recommend appropriate evaluation and treatment. Although they agreed with this plan, the urologists were worried about potential lawsuits, damage to their reputations, and the stress of difficult meetings with the patients and their families.
>
> When the process of notifying the patients started, the hospital president realised that questions about the validity of other biopsies would be raised even though there was no clinical evidence to raise such concern. She thought that all should be reviewed... About 6000 biopsies would have to be reread, and we needed help. Inquiries to the professional pathology bodies were disappointing: not only did we receive little assistance, but we were routinely asked why we wanted to expose more errors..." [683]

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

5.1.2 Mandatory Reporting

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

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

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

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

These regulations illustrate the way in which legal obligations can be placed upon operators so that they are *required* to report certain categories of near-miss incidents. There are examples of similar mandatory systems in other domains. For example, the recent UK National Health Service report entitled 'An Organisation with a Memory' proposed a national mandatory reporting scheme for adverse health events, and specified near misses, based on standardised local reporting systems [635]. There are, however, mixed views about the effectiveness of such systems. For example, the Safe Medical Devices Act of 1990 requires that healthcare facilities and manufacturers must report serious injury or illness related to the failure or misuse of specific medical devices. However, Cohen has argued that:

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

Many of Cohen's criticisms seem to focus on ways in which mandatory systems have been used, or 'abused', by those who operate them. Very few of his adverse comments directly stem from weaknesses in mandatory systems. There are, however, strong concerns about the enforcement of mandatory systems. Clearly, if an individual or group have suppressed information about an incident then others within the organisation must be in a position to detect it if any form of action is to be taken. If an individual fails to report a mandatory occurrence then they run the risk that one of their colleagues may also detect and submit information about an incident. Follow-up investigations might then centre on the reasons why the first operator failed to provide any notification of the adverse event. Alternatively, incidents can come to light through the post hoc review of logs and records. This approach relies upon techniques that are similar to the exhaustive analysis that has been used to identify background incident rates, and thereby derive reporting quotas. Irrespective of the source of such information, there remains the problem of determining what disciplinary action should be taken when individuals fail to report mandatory incidents. Typically, this depends upon the severity of the incident being considered and upon whether the individual had a clear appreciation of that severity. For instance, if the incident occurred during a period of high workload, it may not be certain that the operator did actually detect the adverse event. Even if they did detect it, the pressure of other duties may have prevented them from reporting it. High workload may even contribute to individuals forgetting about low-criticality occurrences [864].

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

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

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

5.1.3 Special Initiatives

Previous sections have argued that voluntary reporting systems suffer from considerable problems of under-reporting. Mandatory systems can address some of these problems, however, they can alienate some members of staff and have not been universally successful. Special initiatives provide an alternate incident reporting technique that can be used to address under-reporting. At their simplest, these initiatives may simply be implemented through simple questionnaires that directly poll staff about incidents and issues that have occurred to them in recent months. This approach has the benefit that all staff may be called on to participate at the same time and in a confidential manner through the return of a simple form. For instance, Sexton, Thomas and Helmreich [736] have exploited this approach to examine more general attitudes to error within the medical profession. Their study prompted returns from 851 operating room staff and 182 intensive care workers. The data that can be obtained using such questionnaires is very different from the more focused information that is provided by conventional incident reporting systems. However, it would have taken many years to elicit the same number of response through more conventional incident reporting systems. More importantly such initiatives can be used to examine the reasons why particular groups *fail* to participate in incident reporting systems even though they may acknowledge that these systems form a valuable part of any safety system:

> "Over 94% of intensive care staff disagreed with the statement 'Errors committed during patient management are not important, as long as the patient improves'. A further 90% believed that 'a confidential reporting system that documents medical errors is important for patient safety'. Over 80% of intensive care staff reported that the culture in their unit makes it easy to ask questions when there is something they don't understand (this is undoubtedly related to the high endorsement of flat management hierarchies in the unit). One out of three intensive care respondents did not acknowledge that they make errors. Over half report that decision making should include more team member input.
>
> More than half of the respondents reported that they find it difficult to discuss mistakes, and several barriers to discussing error were acknowledged. The 182 staff in intensive care reported that many errors are neither acknowledged nor discussed because of personal reputation (76%), the threat of malpractice suits (71%), high expectations of the patients' family or society (68%), and possible disciplinary actions by licensing boards (64%), threat to job security (63%), and expectations or egos of other team members (61% and 60%). The most common recommendation for improving patient safety in the intensive care unit was to acquire more staff to handle the present workload, whereas the most common recommendation in the operating theatre was to improve communication." [736]

As mentioned, questionnaire based techniques are qualitatively different from other forms of incident reporting. They help to reveal general attitudes rather than specific information about particular adverse occurrences. On the other hand, such initiatives are deeply revealing about the attitudes to error that chapter 2.3 has argued to be significant causes of more 'systemic' failures.

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

> Data on side effects of newly launched drugs are limited,1 highlighting the need for effective post-marketing surveillance. An inverted black triangle on product literature identifies new products. Suspected adverse reactions to these drugs, however minor, should be reported to the Committee on Safety of Medicines through the yellow card scheme. Adverse reactions are Adverse reactions are underreported, and few doctors in the United Kingdom know the meaning of the 'black triangle' symbol. We assessed the degree of underreporting of suspected adverse reactions to new drugs in general practice and determined if reporting varied when reactions were severe or previously unrecognised.
>
> There were 3045 events (in 2034 patients) reported as suspected adverse reactions on the green forms during the 10 studies. General practitioners indicated that they had reported 275 (9.0%; 95% confidence interval 8.0% to 10.0%) of these reactions to the Committee on Safety of Medicines: reporting was highest for serious unlabelled reactions (26/81; 32.1 %) and lowest for non-serious labelled reactions (94/1443; 6.5 %). Serious unlabelled and non-serious unlabelled reactions were significantly more likely to be reported than were non-serious labelled reactions. According to general practitioners' responses, the proportion of serious labelled reactions also reported on yellow cards (7/64; 10.9%) was only slightly greater than that of non-serious labelled reactions." [523]

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

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

Issue based incident reporting systems can also be used to make inferences about the background rate of contributions. The difference between the reporting frequency before the initiative and after the initiative can be used both to gauge the success of any focus on particular issues and to provide a more general measure of under-reporting. For example, the UK Meteorological Office operated the European Turbulent Wake Monitoring Scheme between 1995-1999. This was set up by the European Commission because the separation minima at airports take no account of existing meteorological conditions. They are simply calculated using the weight of the aircraft. Under favourable meteorological conditions, however, it may be possible to reduce these minima if a wake vortex is less likely to occur. The intention was to create and maintain a database of wake vortex incident reports with associated meteorological data. Researchers and aviation companies could then use this data to better understand wake vortex behaviour. A voluntary incident reporting system was chosen as a means of compiling this data because fully equipped meteorological monitoring systems cost up to $1m for a single airport. The initiative was intended to address under-reporting problems because the UK was the only European country to regularly monitor wake vortex incidents. However, over 90% of reported encounters in this existing system took place around Heathrow airport. There was "clearly a need for data from airports with a diverse range of runway configurations, meteorological phenomena and capacity in order to assess the global problem" [548]. It can, however, be argued that this scheme illustrates some of the limitations of such focused initiatives. The system ceased to record further data once the initial funding from the European Commission had run out.

5.2 Encouraging the Detection of Incidents

Previous sections have argued that under-reporting continues to be a significant problem for many incident reporting systems. Mandatory participation provides a potential solution but also raises further pragmatic and ethical problems. Special reporting initiatives can be used to assess the scope and nature of the under-reporting problem. However, pro-active questionnaires and systems that are focused on specific types of incidents suffer from different forms of reporting bias. It can also be difficult to sustain high levels of participation in special reporting initiatives. The following pages reviews manual and automated techniques that can be used to combating the problem of under-reporting. These techniques are intended to support the more general class of voluntary incident reporting systems introduced in Chapter 3.7, rather than special purpose or mandatory systems.

5.2.1 Automated Detection

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

Trust and Acceptance

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

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

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

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

As mentioned, the degree of sophistication in the automation that is available to detect potential incidents varies widely from industry to industry. The development of this technology depends both upon the complexity of the application that is being controlled. For example the ability to monitor pilot actions might be interpreted as a by-product of the development more advanced control systems. The development of automated detection technology also depends upon the consequences of failure and the severity of the perceived threat. Although not directly a safety concern, this can be illustrated by recent initiatives to improve the monitoring of security incidents involving US Department of Defence Computers. These represent instances of the malicious failures described in Chapter 1.3:

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

> information on new or improved technology and programs planned for the next 1 to 5 years." [762]

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

Trust and Acceptance

The importance of staff acceptance of automated monitoring devices cannot be underemphasised. The action of trades unions and other forms of worker representation can block the introduction and use of this technology for many years. Driver monitoring systems on UK railways provide a good illustration of this point. In 1999, Her Majesty's Railway Inspectorate (HMRI) issued a report that analysed the management systems which were intended to record and assess incidents in which drivers had passed signals 'at danger'. The number of these incidents that were reported in the UK gradually fell during the 1990's. However, it has levelled out in recent years: 944 in 1991/92, 593 in 1997/98, 643 in 1998/99 These incidents have also led to a number of high-profile accidents. The collision at Watford South Junction on 8 August 1996 caused the death of one passenger. The accident at Southall on 19 September 1997 was also caused by a signal being passed at danger. As a result of these accidents, plans were developed for the introduction of the Train Protection and Warning System (TPWS). This is intended to mitigate the effect of such incidents by warning the driver and ultimately by braking the train at junctions, single lines and 'unusual' train movements, see Chapter 2.3. However, the cost and complexity of such equipment has delayed its introduction. As an interim measure a range of Driver's Reminder Appliances (DRA) have now been fitted to most driving cabs. These have the limited role of reminding a driver of the current signal when they are stopped at a station with the starting signal at danger. Without more advanced protection systems, an argument was made for more closely monitoring driver behaviour. This was based on the idea that human factors problems could be addressed through remedial training and supervision if it was possible to identify those drivers who were most likely to pass signals at danger. The HMRI report reviewed piecemeal progress towards the introduction of driver monitoring equipment that was intended to make this possible:

> "Recommendation 9 of the HMRI report into the accident at Watford South Junction was to North London Railways (now Silverlink) to extend the use of on train data recorders to monitor driving technique. Although the number of trains fitted with the equipment is still less than 20% of the total, the number is increasing rapidly so other Train Operating Companys should be making use of it for unobtrusively monitoring driver performance. Thameslink was not doing so (although was to start) and neither was Connex South Central, although it is acknowledged that most of their fleet is not fitted with the equipment." [351]

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

> "Evidence (of driver involvement in 'signal passed at danger' incidents) should include that to be provided by CIRAS and from On-Train Data Recorders used to monitor driver behaviour. ASLEF in particular should give their full support to such an initiative." [319]

This comment about the need for ASLEF support is important because it reveals the HMRI's sensitivity to workers' concerns about the introduction of these automated sensing systems into the cabs. It is important to emphasise that these concerns again rest on a justified fear of retribution that affects all forms of incident reporting. These fears are exacerbated by a number of additional factors. As we shall see, the sensitivity of these devices can lead to false readings that might, in turn, trigger unwarranted accusations of poor performance and error. The piecemeal and delayed introduction of these systems may mean lead to inequitable treatment. Some driver errors may be 'caught' by these systems whilst others may go undetected because no equipment is installed. Automated equipment triggers the detection of some errors, however, it often fails to capture the 'mitigating' factors that can excuse certain violations. Finally, such automated systems address the observed consequences of deeper systemic failures, including poor signal placement, that actually *cause* the failure in the first place [733] Many of these concerns do not stem from the ethical involved in introducing automated monitoring equipment. Instead, they centre on the ways in management will use the data that is collected by these systems. Such concerns were touched upon by an earlier enquiry. The following quotation is revealing because it probes the limits of a 'no blame' culture. The ambivalent position of the regulators again illustrates how pragmatics lead to what we have termed a 'proportionate blame' approach. The report does not criticise the use of SPAD reports, from either manual or automated sources, within a company's disciplinary procedures.

> "All Train Operating Companys (TOCs) visited have specially monitored driver procedures in place to assign drivers to categories dependent on their incident history. However, the application of these procedures (mandated by Group Standard GO/RT3251) varies widely between TOCs. For example Connex South Central at some drivers' depots allocated all drivers to one of the 'at risk' categories. West Anglia Great Northern Railway's (WAGN) procedure appeared to give rise to too much scope for management discretion in reducing the status of a driver from 'incident prone' to 'normal'. Since, HMRI's inspection, WAGN are revising their procedure.
>
> Drivers in higher risk categories are intended to be subject to a greater number of assessment rides focusing on identified weaknesses, but it is questionable whether these are always achieved in practice. Generally, these are managed by individual drivers' depots, but it would be more satisfactory for this to be monitored centrally within TOCs to ensure that the extra assessments are actually carried out and that they address any identified weaknesses in competence.
>
> There must be adequate procedures for removing a driver from driving duties in the event of their SPAD record not improving despite further training and assessment. Some TOCs use the disciplinary procedure, but the key requirement for TOC managements following incidents is to ensure any deficiencies in competence are identified and robustly addressed by means of further training, if necessary, and competence assessment..." [351]

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

> "The version of GO/RT3252 in use at the time of HMRI's inspection required that when a driver had had three SPAD incidents, they were only to continue on driving duties if there was a written justification for doing so. This was not always found to be the case. Some TOCs were found to take a relatively hard line and removed any driver automatically from driving duties at the third SPAD incident, whereas others did not. It could be significant that those TOCs which were found to take a hard line in this area appeared to have better SPAD records than others, and this may lead drivers to adopt the required defensive driving approach. The new version of GO/RT3252, revised since HMRI's inspection, focuses more on the identification and rectification of competence weaknesses which lead to SPAD incidents. " [351]

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

Missed Incidents and False Alarms

Chapter 2.3 argued that environmental features can prevent operators from accurately perceiving important properties of their environment. Table 5.2.1 was used to show how a signal may or may not be present. If the signal is present then either the operator may detect it, in which case they have achieved a 'hit', or they may fail to detect the signal, this results in a 'miss' in Table 5.2.1. If the signal is absent and the user detects it then this results in a false alarm. However, if they do not detect a signal then this represents a correct rejection. These same distinctions apply both to

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

Table 5.1: Outcomes from Signal Detection

the human detection of signals or warnings in their environment and to the automated detection of critical incidents. For instance, if an automated system detects a signal, that is to say an incident, when none is present then this will generate a false alarm. Conversely, if an incident did occur and was detected then this represents a 'hit' by the detection equipment. A 'miss' occurs if an incident took place but was not detected. A correct rejection takes place when the system successfully finds that no incident has occurred. Wiener summarises the technical problems that emerge from this analysis:

> "In any warning system, one can expect false alarms and missed critical signals, and the designer must design the filter logic to strike a balance. If the system is deigned to be 'sensitive', that is to have a high detection rate, then it will hive a high false alarm rate, and vice versa. There is no perfect system that can detect all true events and filter out all false events." [865]

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

> "...we received two TCAS II-advisories, corresponding to departures. The departures are cleared to 10,000 feet, [and] arrivals...[at] 11,000 feet. The TCAS II reacted to the closure rate of the departing aircraft and our inbound flight. [The] RA was ignored as traffic was in sight. The real problem is that the TCAS II alert caused such a distraction in the cockpit that two or more radio calls from Approach Control were missed." [547]

The conditions that lead to spurious alarms are hard to anticipate. For example, some relate to technical failures in the manner in which aircraft altitude data is acquired from the Mode C function of the aircraft's radar transponder. Should Mode C even temporarily provide erroneous altitude information, an erroneous TCAS II Resolution Advisory command to climb or descend may result [547]. Other false alarms can be generated by local features. For instance, Billings cites numerous spurious warnings at particular airports including Orange County California and Dallas Fort-Worth Texas. He argues that such missed incidents and false alarms have a considerable impact upon the behaviour of system operators. Early versions of the Ground Proximity Warning System (GPWS) were so unreliable that crews ignored or disabled them. Such actions indirectly led to accidents at Kaysville Utah (1977) and Pensacola, Florida (1978) [82]. One large commercial airline discovered 247 (73%) spurious alarms amongst a total of 339 GPWS alerts in a twelve month period.

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

> "Air carrier (X) was inbound on the...STAR level at 10,000 feet. Under my control, air carrier (Y) departed...on the...SID, climbing to [an] assigned altitude of 9,000 feet. Approximately 14 miles SW...I issued traffic to air carrier (X) that air carrier (Y) was leveling at 9,000. Air carrier (X) responded after a few seconds that they were descending. I again told air carrier (X) to maintain 10,000 feet. Air carrier (X) responded 'OK, we've got an alert saying go down'. Simultaneously, air carrier (Y) was getting an alert to climb. They both followed the TCAS II advisorys and almost collided. Later, [the pilot of air carrier (X)] ...indicated [that] his TCAS II was showing zero separation. They passed in the clouds without seeing each other. When pilots start taking evasive action, our equipment cannot update quickly enough for the controller to help. Both aircraft were issued traffic as prescribed by our handbook (merging target procedures). [Air Carrier] Company directives, I'm told, dictate that pilots must respond/follow the TCAS II alert advisories." (ACN 224796) [547]

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

Limited Views of Causation

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

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

> but had him on TCAS II. This seems to be happening more and more...It appears [that pilots]...are using TCAS II instead of looking out the window. As an air traffic controller I cannot have pilots using TCAS for visual separation to maintain spacing (as on one occurrence a crew offered to do). There is no TCAS II separation." ([547], ACN 202301)

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

The particular consequences for incident reporting are that the (ab)use of automated detection systems makes it less likely that personnel will explore the underlying causes of the alarm that they have experienced. This is important because technologies, such as TCASII, minimum safe altitude warning (MSAW) systems, short-term conflict alert (STCA), can be used to trigger investigations that stand some chance of uncovering the deeper systemic issues that exposed users' to danger in the first place. Both Perrow [677] and Reason [702] warn if these underlying causes are not addressed then it is likely that our defences will fail at some point in the future. Each TCAS warning in aviation or SPAD in the rail industry not only warns the individual pilot or driver, it should also be a warning to the industry as a whole.

A number of pragmatic issues limit the amount of information that can be obtained from automated incident detection systems. For instance, TCAS II provides limited data about aircraft separation. It does not provide a 'complete' account of the causal factors and influences that led to the loss of separation. Automated recording equipment can provide more detailed insights into the course of an incident. For example, digital flight data recorders provide information about a failure to fly a stabilised approach, about engine overspeed and about an excessive rate of descent. Over time such data can be collated to provide an overview of common problems, for example repeated overspeeds by several pilots when landing at a particular airport [344]. However, it is important to acknowledge the limitations of the data that can currently be captured. For instance, it is not possible to use digital flight data to determine what caused the specific incidents that are recorded. It may be due to pilot 'error', to air traffic control restrictions, to adverse meteorological conditions etc. In other words, the information that is elicited by automated systems currently only acts as a trigger for further investigation. In this respect, it is no different from the trigger that is provided by the manual detection and contribution of incident reports.

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

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

> the privacy concerns that the pilot associations and others have expressed with respect to recording images of flight crews. As you know, the Board's reauthorisation passed by this Chamber would require that the same protections already in place for CVRs be extended to image recorders in all modes of transportation. Under those provisions, a cockpit image recording could never be publicly released. Even where monitoring has been allowed there is resentment towards certain technologies." [304]

This quotation does acknowledge the concerns that commercial airline pilots feel about the introduction of such systems. These concerns were intensified when several media organisations broadcast the final minutes of the cockpit voice recorder during the Cali crash. Although this would not have been allowed under US or Canadian legislation, there was no provision to prevent the release of such material in Columbia at the time of the accident. There is also the perception amongst pilots that such video equipment is being introduced to satisfy public perceptions about the utility of recordings and that these perceptions may not, in fact, be justified. This argument has considerable strength. Chapter 2.3 noted the difficulty of interpreting intention and cause from video recordings of users who commit 'everyday', non-safety critical errors. These difficulties would be considerably greater in the aftermath of an accident.

Fortunately, near-miss incidents offer alternative means of eliciting additional information to support the output of automated monitoring equipment. Several major airlines have now installed Air Data Acquisition Systems (ADAS) that record a range of information that is typically already recorded by the digital flight data recorders ('black box' recorders). For example, British Airways currently supplement their air safety reporting programme with data collected from their SESMA flight data recorders [660, 661]. ADAS information can be routinely monitored to detect whether certain triggering conditions occur during otherwise normal operation. These triggering conditions include warnings from GPWS, TCAS, stall protection systems etc. They can include attitude transgressions, such as overbanks, or the transgression of speed limits, such as flap overspeeds. They may also include incidents involving extreme g-loads or prolonged flares. If a trigger occurs then the airlines' flight data analysts may interview the crew. Klampfer and Grote used this technique to analyse 71 incidents within a commercial fleet [445]. 48 of the incidents involved A320 aircraft, 18 involved the MD11 and the rest were from a variety of other aircraft. This data revealed that 29% of incidents involved speed violations, analysts included underspeed and overspeed conditions in this category. 19% of incidents involved unstable approaches. 11% involved prolonged flares. 10% involved low go arounds. 10% of all incidents were triggered by the automated monitoring systems mentioned in previous sections. 10% of the incidents involved attitude violations. 8% involved excess g-loads. 3% of the events could not be classified according to these general categories/ The interview data was examined together with the triggering information from the ADAS system. A causal analysis identified that direct human errors contributed to 40% of all incidents. These errors included poor situation awareness and a lack of crew coordination. Human influences contributed to 31% of the incidents. These are classified as actions that are not, of themselves, incorrect but which contributed to or exacerbated the consequences of an incident. This is perhaps the most difficult of Klampfer and Grote's categories; it includes mental overload and routine action as contributory causes. Their analysis also identified that 16% of incidents were caused by environmental factors, including air traffic control 'failure'. Only 11% were caused by technical failures, including poor meteorological information. The remaining 2% were unclassified.

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

Chapters 9.3 and Chapter 14.5 will look at conventional tools, including relational databases, and more advanced techniques, such as case based reasoning, that support the off-line monitoring of incident reports. In contrast, this section has focused on systems that support the on-line, or run-time, detection of adverse occurrences. These systems offer a number of important advantages. In particular, they can help operators to avoid a potential incident or mitigate the consequences once an incident has taken place. The same systems can also be used to trigger further causal investigations after an event has occurred. A particular concern is that most regulatory and commercial organisations focus on the former role of automated detection systems. They often miss the opportunity to address the causes of those incidents that are reported by automated detection equipment. As a result, latent weaknesses become embedded in systems that rely upon detection equipment as primary rather than a secondary defence mechanism.

5.2.2 Manual Detection

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

A Reporting Culture

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

- *local champions.* 'Local champions' promote the system and who act as guarantors. They ensure that assurances of anonymity will be preserved in the face of external or managerial pressures. The previous chapters have already cited the role of David Wright in the local clinical system within Edinburgh's Western General hospital. However, similar comments can be made about some of the much larger systems that operate within major companies. For instance, Capt. Mike O'Leary performs a similar function within British Airways' confidential human factors reporting system [661]. A number of incident reporting systems have explicitly recognised the importance of these individuals. For instance, the Royal College of Anaesthetists advocates the identification of a Critical Incident coordinator who is responsible for drawing up and monitoring the operation of the system [716]. The explicit identification of an individual coordinator is a deliberate policy which goes beyond the more usual use of a committee structure within UK healthcare. There is, of course, a danger that the removal of such key individuals will threaten employee confidence and through that may jeopardise the continued operation of the system.

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

  > "The reportable events indicator is a measure of safety performance but more important than the number itself is the severity of the events reported. No target is set for this indicator in case this should discourage reporting. Indeed, within a healthy safety culture, the introduction of a 'blame free' reporting system may well cause an increase in the number of events reported." [708]

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

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

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

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

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

it produced and improved contingency plans for dealing with oil and chemical spills. Development of the station nature trail continued with habitat management and creation of a wildflower meadow. There are over 1,000 species of flowering plants and invertebrates on the trail, more than 100 of which are currently listed in the Somerset Wildlife Trust list of notable species... A Peregrine Falcon nest ledge on the reactor building has been added to the five other raptor nest boxes situated in and around the nature trail." (Location report: Hinkley point B, [708])

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

Providing Feedback

Effective invention to address acknowledged safety concerns provides what is arguably the most persuasive means of encouraging staff to participate in incident reporting schemes. At the highest level, feedback about safety improvements can be provided through staff publications that record the severity of incidents that are reported each year. For example, Table 5.2.2 presents incident statistics published by the UK Atomic Energy Authority [535]. It shows the number of incidents reported at each level of the International Atomic Energy Authority's International Nuclear Event Scale (INES). INES is used to provide means of comparison between the reports that different national systems submit to the IAEA's INES incident database. Incidents at level 1 are simply regarded as anomalys, level 2 is an incident, level 3 is a serious incident, level 4 is an accident with significant off-site risk, level 5 is an accident with off-site risk, level 6 is a serious accident, level 7 is a major accident. The Chernobyl was classified at level 7, while the 1989 incident at the Vandellos nuclear power plant in Spain was classified at level 3. This did not result in an external release of radioactivity, nor was there damage to the reactor core or contamination on site. Fire did, however, damage the plant's safety systems. The IAEA does not provide examples of incidents below level 3; this is left to the prerogative of individual states. The benefit of this style of feedback is that managers and operators

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

Table 5.2: UK AEA Incident Statistics 1996-1999

can compare national or local safety standards against those of other countries. For example, in 1997 the total INES summary produced for the IAEA recorded 16 anomalies at level 1, 15 incidents at level 2, 2 serious incidents at level 3 and no accidents between levels 4 and 7.

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

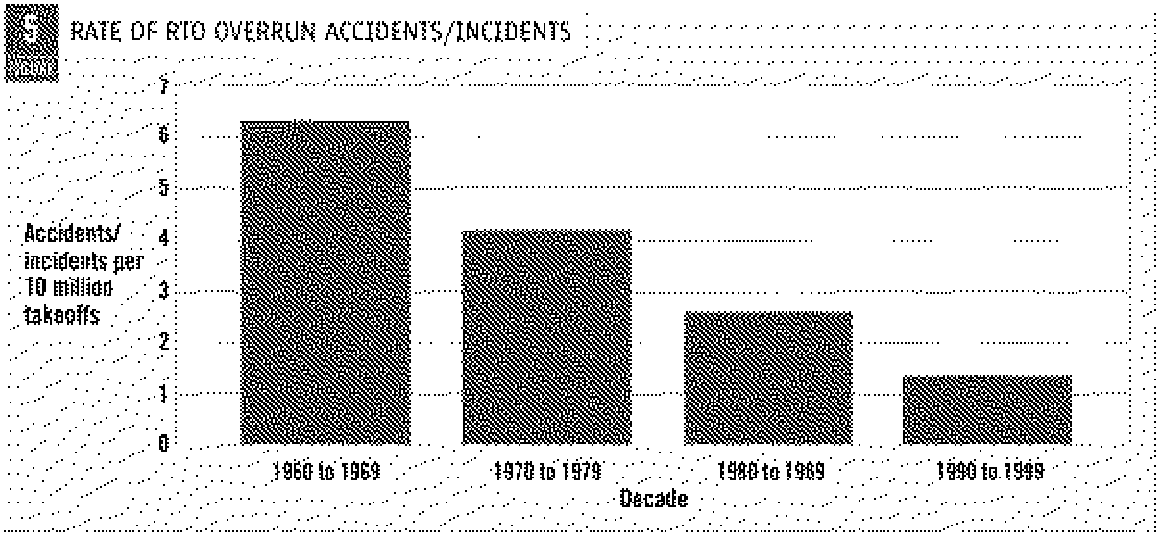

Figure 5.2: Accident and Incident Rates for Rejected Takeoff Overruns

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

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

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

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

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

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

Many reporting systems avoid these criticisms by supplementing raw statistical information with more qualitative accounts and anecdotal editorials about previous incidents. For example, the image on the left of Figure 5.3 illustrates the Feedback reports that are produced by the Confidential Human Factors Incident Reporting Programme (CHIRP) . Feedback is distributed to personnel

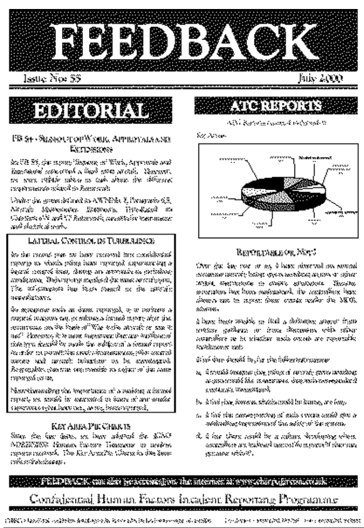

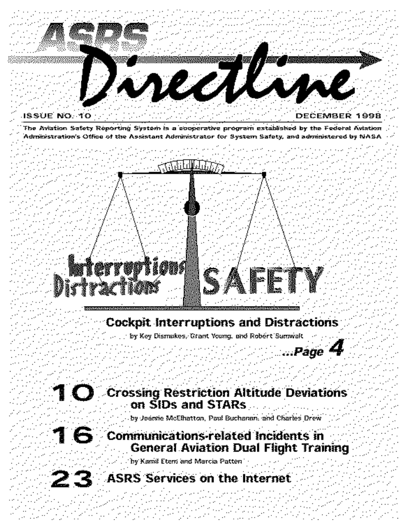

Figure 5.3: CHIRP and ASRS Publications

within commercial and general aviation. They are provided in paper form. They are also available on-line in HTML and PDF formats that can easily be both downloaded and printed. The CHIRP Feedback newsletter has a circulation of approximately 30,000. The ASRS's equivalent publication, Callback, is distributed to 85,000 aviation professionals. As can be seen from the cover in Figure 5.3 statistical data about the frequency and nature of submissions, typically in the form of pie-charts, introduces more qualitative accounts. These are intended to speak 'with the voice' of the individuals who are concerned in an incident:

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

> intermittent, it slipped through and reared its ugly head again the next day during crew checks. It finally reached the point where the line avionics personnel refused to 'shake it up' to get it going, the system needed proper down-time for investigation. Yet on four continuous reports, an A and C engineer with NO avionics clearance or know-how, released the aircraft to service with an inoperative fire detection system. This engineer was a mid-level manager with both a cavalier attitude to anything non-mechanical and also under pressure from management above him. What steps are being taken to address management's limitations to release aircraft to service?
>
> *Editorial comment: The alleged circumstances relating to the release of the aircraft were investigated by CAA (SRG) and corrective actions agreed. In the case of a repetitive defect that has not been cleared after three attempts, the procedure requires that the aircraft be withdrawn from service until the defect is rectified.*" [178]

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

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

> "Why do activities as routine as conversation sometimes interfere with monitoring or controlling the aircraft? Cognitive research indicates that people are able to perform two tasks concurrently only in limited circumstances, even if they are skillful in performing each task separately. Broadly speaking, humans have two cognitive systems with which they perform tasks; one involves conscious control, the other is an automatic system that operates largely outside of conscious control. The conscious system is slow and effortful, and it basically performs one operation at a time, in sequence. Learning a new task typically requires conscious processing, which is why learning to drive a car or fly an airplane at first seems overwhelming: the multiple demands of the task exceed conscious capacity. Automated cognitive processes develop as we acquire skill; these processes are specific to each task, they operate rapidly and fluidly, and they require little effort or attention." [214]

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

Callback, Feedback and DirectLine help to elicit further contributions by explaining how previous incidents can be avoided in the future. These publications are all accessible in electronic form, over the Internet. However, they all rely upon a traditional format. These publications exploit the linear style of conventional newsletters or journals. This has important benefits. In particular, they can be easily printed for wider dissemination. However, a number of incident investigation authorities are exploiting alternative approaches. Most of these are based around providing Internet access to databases of previous incidents. This approach is partly exploited by the ASRS . Reports are published incrementally so that the fifty most recent contributions are summarised in each batch. However, other organisations extend this database approach to include not simply summaries of the incident but also information about the associated investigation and analysis. For instance, the interface on the left of Figure 5.4 provides access to the US Chemical Safety and Hazard Investigation Board's incident reports 'Centre' [162]. This provides access to reports on both accidents, involving

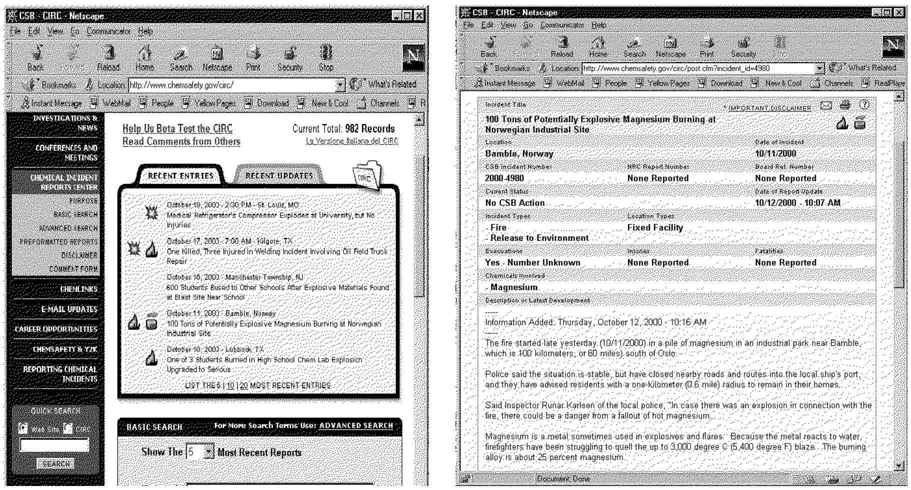

Figure 5.4: Web Interface to the CHSIB Incident Collection

fatalities, and incidents. Users can search through an archive of incident reports using a number of different tools. The image on the right of Figure 5.4 illustrates the information that is returned about each incident. There are a number of innovative features about this application that encourage contributions about adverse occurrences. In particular, it is possible to chart the course of an investigation as it progresses. This provides individual contributors with confirmation that their reports are being acted upon. Search facilities also enable potential contributors to determine whether other similar occurrences have been notified to the system. Chapter 13.5 will provide a more detailed analysis of the dissemination techniques that can be recruited both to publicise the findings of incident investigations and to elicit further contributions to reporting systems. In contrast, this section continues to examine other means of encouraging the manual submission of information about adverse occurrences.

Publicising Procedures and Scoping the System

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

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

> A Critical Incident may be defined as an event which led to harm, or could have led to harm, if it had been allowed to progress. It should be preventable by a change of practice. Complications are occurrences of patient harm, and are sometimes the outcome of critical incidents. If you experience what you think is a critical incident whether or not it has such an adverse outcome, please fill in this form as soon as possible after the event - memory changes very rapidly." [716]

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

> "Please grade how PREVENTABLE the incident or complication was as follows:
>
> 1. Probably preventable within current resource (e.g. failure to do preop machine check);
> 2. Probably preventable with reasonable extra resource (e.g. failure to detect oesophageal intubation would be improved by having capnographs);
> 3. Possibly preventable within current resource (e.g. pneumothorax during CVP insertion might be prevented by better teaching and supervision);
> 4. Possibly preventable with reasonable extra resource (e.g. problem arising because anaesthetist unwell might be prevented by more cover);
> 5. Not obviously preventable by any change in practice (e.g. electricity grid failure)" [716]

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

> "The reporting criteria for CAIRS injury/illness cases changed, effective January 1, 1990, from the criteria specified in the DOE Guide to the Classification of Recordable Accidents to the Occupational Safety and Health Administration (OSHA) guidelines. The reporting threshold for property damage cases was originally set at $1,000 and remained that way until January 1, 1996, when it was raised to $5,000. The vehicle accident reporting threshold was $250 from 1975 through 1985, $500 from 1986 through 1995, and was raised to $1,000 effective January 1, 1996." [657]

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

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

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

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

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

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

5.3 Form Contents

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

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

5.3.1 Sample Incident Reporting Forms

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

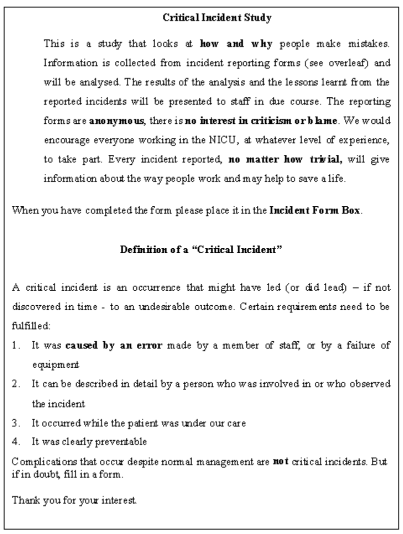

Critical Incident Study

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

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

Definition of a "Critical Incident"

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

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

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

Thank you for your interest.

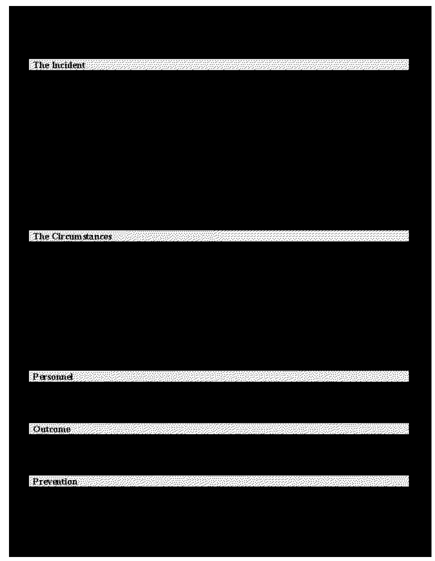

Figure 5.5: Incident Reporting Form for a UK Neonatal Intensive Care Unit [119]

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

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

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

Table 5.3: Developing Reporting Forms

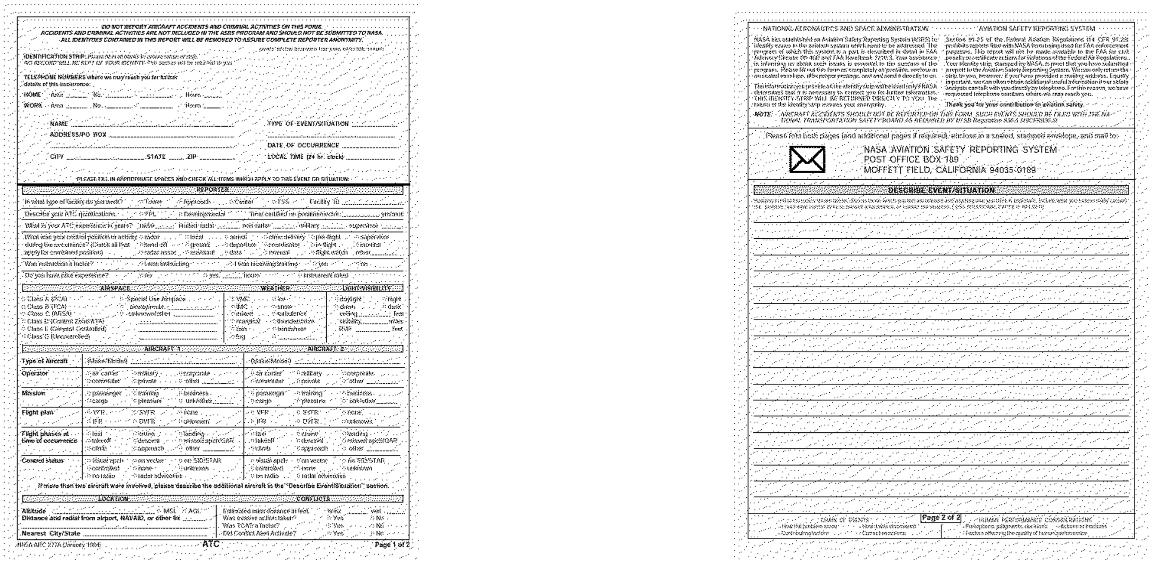

Figure 5.6: ASRS Reporting Form for Air Traffic Control Incidents (January 2000)

ASRS report forms cannot, however, be submitted using Internet based technology. There are clear problems associated with the validation of such electronic submissions. A small but increasing number of reporting systems have taken this additional step towards the use of the Web as a means both of disseminating and submitting incident reporting forms. For instance, Figure 5.7 illustrates part of the on-line system that has been developed to support incident reporting within Swiss Departments of Anaesthesia [757].

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

5.3.2 Providing Information to the Respondents

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

Assurances of Anonymity or Confidentiality

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

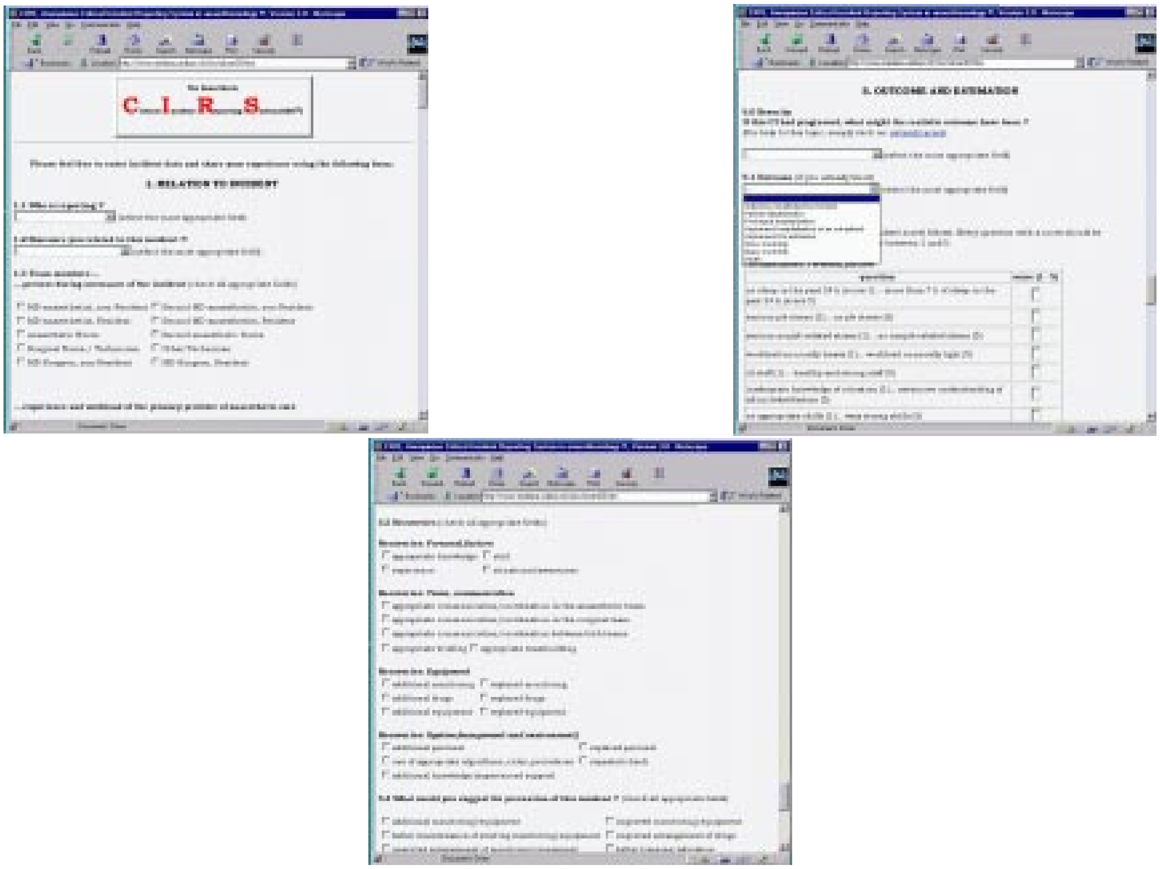

Figure 5.7: The CIRS Reporting System [757]

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

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

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

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

The Swiss scheme shown in Figure 5.7 states that:

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

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

Definitions of an Incident?

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

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

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

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

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

It is worth considering the implications of this definition in the light of previous research in the field of human error. For example, Reason has argued that many operators spend considerable amounts of time interacting in what might be terms a 'sub-optimal' manner [700]. Much of this behaviour could, if left unchecked, result in an undesirable outcome. However, operating practices and procedures help to ensure safe and successful operation. From this it follows that if respondents followed the literal interpretation of the CIRS definition then there could be a very high number of submissions. Some schemes take this broader approach one step further by requiring that operators complete an incident form after every procedure even if they only indicate that there had been no incident. The second interesting area in the CIRS definition is the focus on team working. The number of submissions to a reporting system is likely to fall as the initial enthusiasm declines. One means of countering this is to launch special reporting initiatives. For instance, by encouraging users to submit reports on particular issues such as team co-ordination problems. There is, however, the danger that this will lead to spurious attention being placed on relatively unimportant issues if the initiatives are not well considered.

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

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

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

Explanations of Feedback and Analysis

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

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

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

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

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

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

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

5.3.3 Eliciting Information from Respondents

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

Detection Factors

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

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

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

Causal factors

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

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

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

Consequences

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

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

Mitigating factors

Several authors argue that more attention needs to be paid to the factors that reduce or avoid the negative consequences of an incident [842]. These factors are not explicitly considered by most reporting systems. There is an understandable focus on avoiding the precursors to an incident rather than mitigating its potential consequences. For instance, the ASRS forms simply ask respondents to consider "Corrective Actions" as a footnote to the free text area of the form shown in Figure 5.7.

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

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

Prevention

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

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

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

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

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

5.4 Summary

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

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

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

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

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

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