

Chapter 10

Causal Analysis

This book is based around an implicit model of incident reporting. The evidence that is collected during primary and secondary investigation helps to reconstruct the event leading to an adverse occurrence. The resulting models and simulations can then be analysed to distinguish root causes from contributory factors and contextual details. Previous chapters have briefly introduced the analytical techniques that can be used to identify the most salient events from a more general reconstruction. The following pages build on this by describing the aims and objectives of such techniques in more detail.

10.1 Introduction

Chapters 7.3 and 8.3 have described how simulation and modelling techniques can be used to reconstruct the events that lead to failure. Causal analysis looks beyond *what* happened to identify the reasons *why* [250]. Kjellén [444] identifies three broad approaches to causal analysis:

- *Expert judgement.* Even with the support of analytical and statistical techniques, mentioned below, it is difficult to prevent investigators from forming subjective judgements that help to shape and direct the causal analysis of any incident. These judgements influence every stage of the investigatory process and so can have a profound impact upon the nature and extent of the evidence that is obtained before any causal analysis even begins. It is important to emphasise that subjective judgements need not be ‘a bad thing’. They reflect the expertise and experience of the investigator. As we shall see, many of the recommended analytical and statistical techniques do little more than document the process of forming these judgements so that they are open to challenge, or validation, through peer review;
- *Statistical techniques.* These techniques are, typically, applied to identify common causal factors amongst a collection of similar incidents. They help to determine whether the presence or absence of certain factors increases the probability that an incident will occur. At its simplest, statistical techniques can track uni and bi-variate distributions [471]. Chapter 2.3 has, however, argued that many incidents have complex, inter-connected causes. Chapter 4.3 has also argued that the limitations of automated logging systems and the unreliability of human witnesses can filter the evidence that is obtained about these causal factors. Some researchers have, therefore, begun to explore multi-variate techniques. There have also been some initial attempts to exploit Bayes’ theorem as a means of quantifying the likelihood that a particular root cause led to an incident given particular observations of that incident. This work builds upon attempts to assess the reliability of software systems given uncertain information about potential failure modes [497]. These techniques can also be used post hoc to parameterise expert assessments about the likely causes of an incident. Chapter 14.5 will describe this work in more detail. It will also describe the practical limitations that have restricted the application of these more advanced statistical techniques;

- *Analytical techniques.* These techniques provide a broad range of formal and semi-formal techniques that are intended to support causal analysis. Many of these approaches rely upon counterfactual reasoning. Causal factors can be distinguished from contributory factors and contextual details if it can be argued that if the causal factor had not occurred then the incident would not have occurred. As we shall see, it can be difficult to apply this form of reasoning to certain incidents especially when the failure of a barrier is identified as a potential causal factor. Other analytical techniques, therefore, rely upon checklists. Investigators are guided in their causal analysis by a limited number of pre-defined categories that help to identify common factors in previous incidents. These approaches are limited in that investigators may be biased towards particular categories, for example those that appear at the top of a list. This can hinder a more coherent causal analysis.

Chapter 10.4 surveys a number of different statistical and analytical techniques. Chapter 14.5 also provides an overview of the use of statistical techniques in monitoring the changing causes of incidents between industries and within the different groups of similar organisations. In contrast, this chapter focusses on the detailed application of one particular set of analytic tools. The strength and weaknesses of techniques advocated by NASA [572] and by the US Department of Energy [210, 209] are demonstrated using the loss of the Mars Climate Orbiter and Mars Polar Lander. Neither of these case studies is ‘safety-critical’. They have, however, been chosen because they illustrate the general applicability of incident reporting techniques to investigate the failure of dependable systems. These two concepts are closely related [851]. Their similarities can be used to advantage of borrowing techniques from one to deal with the other [486]. The NASA case studies were also chosen because of the technological sophistication of the systems involved, they therefore represent a strong contrast with the National Transportation Safety Board (NTSB)’s Allentown incident in Chapter 8.3.

It can, in practice, be difficult to distinguish between the stages of investigation, reconstruction and analysis. Investigators may be forced to obtain more evidence to resolve the omissions and ambiguities that are identified when they reconstruct the events leading to failure [462]. Similarly, investigators often have to extend the scope a reconstruction as new theories are developed about the cause of an incident. Chapter 5.4 has also described how the collection of evidence can be biased, or ‘focussed’, by an investigator’s working hypotheses about the probable course of events. These pragmatic issues can complicate the application of the modelling techniques that have been introduced in previous chapters. The costs associated with the development of interactive three-dimensional simulations can dissuade investigators from revising them in the light of new causal hypotheses. Similarly, the problems of maintaining large and complex graphical models can force investigators to use techniques that have stable tool support. The closing sections of this chapter, therefore, attempt to assess the practical implications of the analytical techniques that are introduced. In particular, there is a concern to assess the degree to which these approaches support ‘real world’ investigation practices [73].

10.1.1 Why Bother With Causal Analysis?

Incident analysis techniques, typically, provide means of distinguishing root causes from contributory factors and contextual details. Chapter 6.4 introduced these different causal concepts. They can be summarised as follows. A causal factor was described using a counterfactual argument [491]. If a causal factor had not occurred then the incident would not have occurred. If A and B are states or events, then A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either. It is important to emphasise that this is based on Mackie’s idea of singular causality [508]. Singular causality is used because there may be other failures that could have had the same consequences but which did not occur in this instance. In contrast, root causes depend upon a more general view of causality. These are causes that have the potential to threaten the safety of future systems. They may, in turn, contribute to a number of the causal factors that are observed in a particular incident. In contrast, contributory factors can be thought of as individually necessary but not globally sufficient [678]. These are events or conditions that collectively increase the likelihood of an accident but that would not themselves lead to an adverse

occurrence. Finally, contextual details are events or conditions that did not directly contribute to an incident. They help to set the scene and establish the context in which an adverse occurrence took place. They can also help to establish that certain factors were NOT significant in the events leading to failure.

It might seem superfluous to ask why analytical techniques have been developed to distinguish between the factors described in the previous paragraph. It is clearly important to analyse the circumstances of a near miss to determine how best to avoid any recurrence that might result in more severe consequences. Within this high level goal, there are a number of more detail motivations for incident analysis. These different motivational factors can have an important effect in determining which analytical techniques will offer the greatest benefits for any particular organisations. These justifications for incident analysis can be summarised as follows:

- *analysis is a regulatory requirement.* In many industries, organisations must analyse their incident reports in order to meet regulatory requirements. For example, ICAO Annex 13 requires that member states not only analyse the causes of individual aviation incidents but also that organisations must use this analysis to identify any common causes between similar reports [386]. Similarly, the UK Rail Inspectorate’s assessment criteria for safety cases requires that all operators demonstrate “established adequate arrangements for identifying the causes of incidents” [352]. Even if there is no regulatory requirement, institutional and organisational policy often requires that a causal analysis should be performed. For instance, the US Army has published detailed recommendations that can be used to determine potential causal factors during an incident investigation [804]. NASA have published similar guidelines [572].
- *analysis is a prerequisite for statistical comparisons.* Regulators are concerned to ensure that organisations identify the causes of potential incidents. This is important if companies are to learn from previous failures. Companies must also analyse the causes of potential incidents because regulators use this information to target their intervention in the market place. Causal information from individual companies is, typically, entered into a central database. This database is then queried at regular intervals to identify common causal factors and also to generate a ‘most wanted’ list of safety improvements within an industry. The UK Health and Safety Executive recently announced its initiative reduce the fatality and major injury rate from 260 per 100,000 workers in 1999/2000 to 230 per 100,000 workers by 2009/2010. Together with these targets they have also announced a review of their incident reporting regulations [337]. The HSE recognise that the overall effectiveness of any safety intervention is determined by the regulator’s ability to identify the root causes of common incidents. The review indicates the need to have confidence in the analytical and reporting procedures that inform each statistical return.
- *focus for remedial actions.* The most immediate reason for performing a causal analysis is to focus remedial actions in the aftermath of an incident. Short-term resources should address the root cause before any contributory factors. Once investigators have addressed immediate concerns over the root cause of an incident, additional resources can be allocated to other events and conditions that contributed to the incident. It is apparent, however, that any disagreement about the causes of an adverse occurrence can have profound consequences. Similarly, significant problems can arise if the analysis fails to correctly identify the root cause of an incident. Under such circumstances, the investigators’ ability to prevent a potential recurrence will be compromised by the allocation of resources to less significant aspects of a system. This is illustrated by the way in which poor training is often identified as a root cause of medical incidents rather than the poorly designed equipment and long working hours that staff are forced to endure [121].
- *guiding the allocation of development resources.* At an organisational level, incident reporting schemes are often argued to be an effective means of informing risk analysis. As we shall see, however, many organisations do root cause analysis but do not feed the data into design. Information about previous failures can be used to direct both acquisition and development work. Such an integrated approach can only be successful if organisations can correctly identify

those components and processes that contributed most to an incident. If the analysis of an adverse occurrence is biased by political or organisational pressures then there is a danger that other aspects of a system will be unnecessarily implicated in the causes of an incident. Long-term development resources may be allocated to areas that do not pose the greatest threat to future incidents. This is illustrated by the Fennell report which argues that the London Underground Management "...remained of the view that fires were inevitable in the oldest most extensive underground system in the world" [249]. The root cause of these fires, in particular the built up of detritus in key areas of the system, was not addressed. Instead, staff were trained to detect and respond to these incidents once they had started. There continued to be a steady number of minor fires until the Kings Cross' accident.

- *characterisation of causal complexes.* The causal analysis of incidents need not simply focus on identifying a single root cause. This has been a weakness in the statistical returns that have been required by some regulators. As many authors have observed, incidents and accidents typically stem from pathological combinations of events [700]. As much can be learned from the ways in which those failures combine as can be learned from single causal factors in isolation. This poses a number of problems. Rather than describing safety priorities in terms of a 'hit list' of individual causal factors, it may be more important to identify critical patterns of events. For example, the recruitment of a new sub-contractor followed by a component failure or the installation of a new item of equipment shortly before a software release. It is for this reason that many organisations, including the European Space Agency and the US Navy [5], have begun to look beyond simple categorisations of causal factors. Later sections will describe this 'lessons learned' work in more detail. For now, however, it is sufficient to observe that they have developed data mining and information retrieval techniques that help investigators to identify patterns within a collection of previous incidents [416].

These motivations provide criteria that can be used to assess the utility of different analysis techniques. For example, the previous chapter briefly explained how the minimal cut set of a fault tree can be used to support incident analysis. The elements of this set represent the smallest possible conjunction of events in which if any basic event is removed then the top condition will not occur [27]. Root causes are basic events that are common to every member of the minimal cut set. There is no reason why there should not be multiple root causes that are common to the elements of this set. In consequence, this approach cannot easily be applied to identify a unique root causes.

There are further tensions between the different motivations that support the causal analysis of near miss incidents. As we shall see, some analytical techniques identify a 'primary causal factor'. These techniques, typically, require that investigators select the most significant cause from a predetermined list of potential factors. This approach helps to ensure consistency between different investigators. The use of an agreed list helps investigators to avoid using a range of different terms to describe the same causal factors. This can, in turn, increase confidence in regulatory statistics. There are, however, a range of problems. It can be difficult to construct an appropriate list of agreed causal factors. As we have seen, new causal factors can emerge with the introduction of novel equipment and working practices. It can also be difficult to identify a single 'main' cause from many competing alternatives. Previous sections have shown how a single event can have multiple proximal and distal causes. Any one of these could be regarded as a root cause on the basis of Lewis' counterfactual arguments. For example, the Allentown incident might have been avoided if excess flow valves had been installed or if proper excavation procedures had been followed. Which of these is the true 'primary' cause?

This analysis illustrates a number of points that will be reiterated throughout this chapter. Firstly, analytical techniques must often be refined to support particular organisational objectives. For example, investigators are often expected to translate their findings into a form that is acceptable to regulatory organisations. This can involve the selection of a primary causal factor from an 'accepted' list of root causes. There is a danger that such requirements may prevent investigators from adequately considering the complex causes of many technological failures [677]. Secondly, causal analysis can yield important information for the subsequent development of safety-critical applications. It is, therefore, important that the products of such an analysis should be in a form

that is compatible with subsequent risk assessment procedures. This does not imply that similar techniques should be used for both activities. However, it is important that designers can understand the outcome of any causal analysis. Finally, the term ‘causal analysis’ applies at several different levels. The previous discussion has used it to describe the process by which the root causes of a particular incident can be distinguished from contributory factors and contextual details. However, causal analysis can also be applied over collections of incidents. This is essential if investigators are to identify patterns of failure and emerging trends in a number of similar incidents.

10.1.2 Potential Pitfalls

Previous paragraphs have introduced some of the complexities that affect the causal analysis of adverse incidents. For example, regulatory requirements impose additional constraints upon the causal analysis of some incidents. The format that best supports ‘organisational learning’ may not be the best format to support the statistical analyses demanded by regulators. There are further complexities. In particular, analysts may lack the evidence that is necessary to perform a detailed causal analysis. Later sections will describe how design decisions and budgetary constraints determine that NASA’s Mars Polar lander would not provide any telemetry data during the Entry, Descent and Landing phase of the mission. In consequence, it was impossible for investigators to accurately reconstruct the events that led to the failure nor could they identify definitive root causes. The following paragraphs, therefore, examine further problems that can complicate the analysis of ‘near miss’ incidents:

- *The scope of a reporting system influences the scope of any causal analysis.* In an ideal situation, investigators would conduct an analysis in an environment that is free from external or organisational constraints. Unfortunately, this does not reflect the experience of most operational reporting systems. For example, local schemes deliberately restrict the scope of the investigator’s analysis to ‘target the doable’. Many hospital reporting systems identify failures within a particular department or ward [119]. They explicitly exclude reports that deal with failures in other departments or at higher levels in the management structure. This pragmatism effectively restricts the scope of any analysis to the immediate activities of the group that participates in the reporting scheme. Of course, the scope of any analysis can be widened as reporting systems are extended nationally and across an entire industry. In consequence, national and international reporting systems are being developed within the healthcare industry. However, these initiatives also place either explicit or implicit boundaries on the scope of any investigation. For example, the ASRS was deliberately established to cut across the many different professional and organisational demarcations that characterise the US aviation industry. It solicits input from commercial, military and general pilots. It encourages reports from air traffic controllers and ground staff. It is important to remember, however, that even this scheme is bounded by organisational factors. For instance, the ASRS provides relatively few insights into ‘near miss’ incidents involving military aircraft. This partly stems from a noticeable under-reporting, mentioned in Chapter 4.3. It also arguably reflects the ASRS’ analytical focus on commercial and general aviation.
- *Organisational factors place unnecessary constraints upon causal analysis.* Organisational goals and priorities influence any causal analysis. These influences do not simply act upon the individuals who report adverse occurrences. They must also affect incident investigators. The most obvious manifestation of this is the lack of critical analysis about regulatory intervention. As noted in the opening chapters, regulators are ultimately responsible for the safety record in most industries. Very few investigators ever analyse the impact that these organisations have upon the course of an incident. There are some notable exceptions to this, including the NTSB’s Allentown report that was cited in the previous chapter [589]. These exceptions, typically, occur when investigators are independent both from the regulator and from any organisation that is directly implicated in an incident. In particular, regulatory failure is most often exposed at the large scale public enquiries that follow major accidents [194]. Given the

pragmatics of most reporting systems, it should not be surprising that such causal factors are not more apparent in the analysis of ‘near miss’ incidents.

- *Organisational can inform a causal analysis.* The previous paragraphs have stressed the way in which organisational factors can constrain the scope of any causal analysis. It is also important to emphasise that these factors can play a positive role. In particular, the last decade has seen a movement away from individual blame as a satisfactory causal interpretation of many incidents. This movement has been promoted by many researchers [702, 845]. However, their work would have had little weight if commercial and regulatory organisations had not had the insight to act upon it. In particular, it is important not to underestimate the powerful normalising influence that investigator training can have upon the products of any causal analysis. This can be seen in the impact of Crew Resource Management training in the aviation industry. This has equipped investigators with a vocabulary that can be used to describe the causes of failure in team-based communication and decision making. Before the widespread introduction of this training, investigators failed to derive many insights about the role of team factors in the causes of many incidents and accidents [57, 735, 412].
- *Historical factors help to shape any causal analysis.* The previous paragraph has argued that explicit training can inform an investigators’ interpretation of the events leading to an incident. Implicit forms of training also play an important role in determining the outcome of any causal analysis. For instance, traditions of interpretation can become established within groups or teams of investigators. This can be seen as a strength; similar incidents are handled in a consistent manner. There is, however, a danger that investigators will become habituated to some causal factors so that they are identified irrespective of the circumstances surrounding a particular incident. In the past, human error was often seen as a routine cause of many incidents [719]. Increasingly, however, software is being identified as the predominant cause of many safety-critical incidents and accidents [413]. For example, later sections will describe the software failures that led to the loss of NASA’s Mars Climate Orbiter and to difficulties in the Stardust programme. These failures clearly helped to focus the investigators attention on software failure as a potential factor in the subsequent loss of the Mars Polar Lander. It is important that the causes of previous incidents inform rather than bias subsequent investigations. This narrow distinction raises important pragmatic problems for investigators who must retain an open mind when they deploy finite analytical resources.
- *Causal analysis is constrained by available resources.* The second half of this chapter will present a range of analytical techniques that investigators can use to distinguish root causes from contributory factors and contextual details. These approaches differ in terms of the amount of time that investigators must invest before they can learn how to exploit them. They also offer different levels of tool support. These factors can have a profound impact upon which analytical techniques are chosen within a particular organisation. More complex techniques are less likely to be used in local reporting system that must rely upon the enthusiasm of key individuals with limited training in incident analysis. Resource constraint also affect national and regional systems. Investigators must justify resource expenditure to upper levels of management if they are to ensure continued support for a reporting system. This topic is addressed in the final chapters of this book. As we shall see, it is difficult to underestimate the importance of these cost-benefit decisions. Complex techniques will fail to provide analytical insights if they are under-resources. Conversely, these more advanced approaches often carry a significant overhead in terms of staff time that cannot be justified for many relatively simple incidents. However, it is equally important to emphasise that ‘low-cost’ analytical techniques often yield superficial results when they are applied to more complex incidents. The problem of selecting an appropriate analytical technique is compounded by the lack of empirical evidence, or published practical experience, that compares the costs and benefits of different forms of causal analysis.
- *Who Performs the Analysis?* The previous paragraphs provide an insight into the complexities that surround any causal analysis of adverse occurrences. As can be seen, many of these

issues focus upon the organisational biases that affect any investigation. These biases can have both positive and negative influences with respect to the overall safety of an application. For instance, an emphasis away from individual error can be beneficial in encouraging investigators to look for wider causes of adverse occurrences. Similarly, by focusing on the ‘doable’ investigators can maximise the allocation of their finite resources. Organisational factors have a negative impact if individual or group objectives are considered to be more important than the overall safety of an application. It is for this reason that many reporting schemes rely upon outside organisations to analyse the reports that they receive. For example, the University of Strathclyde coordinates the analysis of incident data on UK Railways [198]. The ASRS is operated by Batelle under contract from NASA. These external organisations assume responsibility for the analytical techniques that are then applied to each report. This approach has the benefit that investigators are seen to be independent from the organisations who must act on any recommendations. In practice, however, there remain strong implicit constraints on the forms of analysis that are performed even by external investigators. For example, a semi-competitive tendering process is often used to award the contracts for these systems. This process can focus the attention of the existing contract holder. It can also introduce terms of reference within a contract that place specific bounds on the form of analysis that is to be performed.

- *The Importance of Balancing Domain Expertise and Multi-Modal Skills.* The emergence of national and international systems has seen a new generation of professional incident investigators. These analysts fall into one of two categories. Firstly, domain specialists often ‘move’ into incident investigation after lengthy periods of field service. There are strengths and weaknesses to both approaches. Domain specialists can quickly lose touch with current operating practices in rapidly changing industries. In consequence, they must either undergo continual retraining to reinforce their existing skills or they must gather new ones. In particular, domain specialists often lack expertise in the human factors domain, they may also have little first hand experience of systems engineering. This makes their analysis vulnerable to criticisms from individuals with these more specialist skills. Secondly, there is a growing number of incident investigators who are recruited in spite of their lack of domain skills. These individuals contribute what can be termed ‘multi-modal’ analytical techniques. They provide tools from other engineering disciplines, such as human factors and systems engineering, that can be applied to analyse incidents in many different application domains. The situation is then reversed, the analytical insights provided by these individuals is then vulnerable to criticism by those who have first hand experience of the application domain. Such observations should emphasise the political nature of many investigations; there is a danger that any analysis may be jeopardised by disagreements between domain specialists and expert witnesses who possess these multi-modal skills. Some organisations, notably the Australian Transportation Safety Bureau, have launched a series of initiatives that are intended to find some middle ground [49]. They have deliberately distinguished between multi-modal and industry specific training requirements. Investigators from each mode of transportation are expected to possess these multi-modal skills, including human factors and systems engineering expertise. In addition, they must refresh the technical and practical foundations of their domain knowledge. However, the ATSB intend that their inspectors will be qualified in more than one domain. This will help to transfer multi-modal analytical techniques between road, rail, maritime and aviation investigations. Just as the US NTSB have established a reputation for their innovative use of simulation and reconstruction techniques, the ATSB continue innovate in the way that they train and deploy their investigators. It remains to be seen whether this transition from a narrow focus on domain expertise to a multi-modal approach will have a lasting impact on the nature of incident analysis within each mode of transportation.
- *The Importance of Justifying Causal Analysis.* The mutual vulnerability of domain specialists and multi-modal investigators raises a number of important concerns about the application of analytical techniques within many investigations. In particular, the individual investigator’s interpretation of an incident is open to many different challenges. It is, therefore, very im-

portant that sufficient evidence is provided about the analytical techniques that are used to support the findings of any investigation. This has been a particular weakness of investigations into human factors issues. Frequently investigators refer to problems of high workload and poor situation awareness without explaining the particular observations that support these conclusions [410]. Of course, as noted above, not all of these analyses were performed by investigators with the relevant human factors training. Similar weaknesses can also be found in systems engineering accounts. For example, it is often difficult to replicate the vibrations that metallurgists have identified as a primary cause of metal fatigue in aircraft components. The ambivalent results of airborne and ground tests are occasionally omitted. In other instances, investigators place sparse details of negative results in appendices that are not then distributed with the body of a report. It can be argued that these techniques support the dissemination of important safety information. Most readers are unconcerned with the methods that were used to reach a particular conclusion. However, these same techniques can be viewed as rhetorical devices. The lack of analytical detail prevents other investigators from raising detailed objections to an analysts findings. It is for this reason that I believe all investigators should provide detailed documentation to support the findings of any analytical technique.

- *Avoid the over-interpretation of sparse data.* There are many reasons why investigators must document and justify their use of analytical techniques. In particular, there is a danger that individuals will be tempted to form conclusions that are not warranted by the evidence that is available. This tendency can be exacerbated by some of the factors that have been mentioned in previous paragraphs. For example, limited resources can force investigators to identify causal factors that are characteristic of a class of incidents rather than analyse an incident for any distinguishing characteristics. Alternatively, organisational pressures can persuade investigators that an incident supports some more general political argument. The ambiguous nature of many incidents can make it difficult to resist such influences. As we have seen, adverse occurrences typically have many potential causes. Given sparse data, limited resources and the pressure to act, it is hardly surprising that some investigators are tempted to ‘cut corners’. Such practices often only come to light in the aftermath of a major accident. This is illustrated by the treatment of Signals Passed at Danger (SPADs) on UK railways. Chapter 6.4 quoted the report from Her Majesty’s Railway Inspectorate, which found that “in some cases greater emphasis was placed on completing a multi-page form than getting to the root cause of the SPAD incident” [351]. Incident investigations tended to focus on issues of driver vigilance rather than the placement of signals or on the other protection mechanisms that were intended to prevent these incidents from occurring. The HMRI report concluded, investigators might have looked deeper into these incidents if they had been required to follow more rigorous techniques for root cause analysis.
- *The Problems of Ambiguous and Limited Evidence.* Incident reconstructions help to establish *what* happened. Causal analysis then identifies the reasons *why* an incident took place. As we have seen, however, these distinctions are difficult to maintain during an incident investigation. Causal hypotheses are formed and reformed as new evidence is obtained about the course of an incident. This creates problems because the resource limited nature of many enquiries can force investigators to develop ad hoc stopping rules. These involve procedures to help them decide when to stop gathering more evidence in support of their analysis. Typically, these procedures involve team presentations or discussions with safety management who must then authorise the end of an investigation. Other circumstances can prematurely curtail a causal analysis. For instance, there may be little direct evidence about the events that led to an incident. Paradoxically, however, NASA’s Mars Polar Lander report demonstrates that a lack of evidence does not bring a causal investigation to a premature conclusion [580]. In contrast, it opens up a vast number of possible explanations that must be discounted before reaching a tentative conclusion. In assessing the analytical techniques that will be presented in this chapter, it is therefore important to remember that investigators may have to use them to discount certain hypotheses as well as to support others.

- *The Problems of Intention.* The previous paragraph has argued that causal analyses are complicated by a lack of evidence about the events leading to a failure. This evidence, typically, relates to the observable behaviour of system components. Similar problems are created when analysts lack information about less visible influences on the course of an incident. In particular, it can be difficult to determine the role that human intention plays in an adverse occurrence. Chapter 2.3 has introduced numerous distinctions between different forms of error and violation. In practice, however, investigators often lack the information that is necessary to distinguish between these different forms [869]. For instance, mistakes stem from an inappropriate intention. It can be difficult for individuals to admit to such intentions in the aftermath of a near miss incident. These problems also affect the interpretation of human behaviour captured on video and audio logs. For instance, individuals have been observed to act in bizarre and pathological ways. They have disregarded operating procedures and violated safety requirements through factors as diverse as boredom, curiosity and a sense of fun [864]. It seems apparent that the advocates of cockpit video recorders significantly underestimate the problems of interpreting human intentions from the behaviour that is captured by these devices. Pedrali's video analysis of optimal and sub-optimal behaviour in commercial test pilots provides ample evidence of this [674]. Later section will describe how ethnographic and work-place studies have been proposed as means of supporting the eventual analysis of such behaviours.
- *Inter-Analyst Reliability.* Many of the problems described in this section stem from a meta-level concern that investigators should be able to replicate any analysis of an incident. This is supported if investigators justify their decision to use a particular technique to support their causal analysis. They should also document any intermediate findings that emerge to support or refute particular conclusions. These requirements enable others to replicate the application of particular analytical techniques. They will not, of course, enable others to directly replicate the results of any causal analysis. Lekberg's work has shown that these results are not simply determined by the choice of an analytical technique [484]. They are also determined by the educational background of the investigator. McElroy has provide a preliminary validation of these ideas [530]. His work showed that even when analysts are trained to use one of the more advanced techniques for causal analysis, their findings will vary considerably even for the same incident. Such problems can be addressed by ensuring that the analysis is replicated by a sufficient number of analysts. This form of mass replication can be used to minimise individual differences in interpretation. However, this averaging out can often lead to polarised views within a team of investigators and it is not clear that a consensus must emerge from replicated forms of analysis. In addition, most reporting systems cannot afford to validate their conclusions through the repeated replication of a causal analysis. There can, therefore, be little confidence that any of the techniques in this chapter will ensure inter-analyst reliability. This is true even for techniques that are supported by formal proof techniques; investigators may disagree about the choice of abstractions that are used within a model. Causal reasoning techniques do, however, increase the transparency of any investigation. They help to document the methods that were used to support particular findings about the causes of an adverse occurrence.

The previous paragraphs provide a stark assessment of the many problems that complicate the causal analysis of safety-critical incidents. These range from pragmatic issues of funding and resource management to the more theoretical barriers to interpreting intentions from observations of human behaviour. Later sections in this chapter, therefore, review some of the solutions that have been proposed to address some of these concerns. In contrast, the following pages describe two incidents that are used to illustrate this comparative study of analytical techniques.

10.1.3 Loss of the Mars Climate Orbiter & Polar Lander

In 1993, NASA commissioned a program to survey the planet Mars. The Jet Propulsion Laboratory (JPL) was identified as the lead centre for these missions. Lockheed Martin Astronautics was selected

as the prime contractor. The program initially consisted of the Mars Global Surveyor (MGS), to be launched late in 1996. This global mapping mission is currently orbiting Mars. The Mars Surveyor'98 project was intended to build on the Global Surveyor's work. This program consisted of the Mars Climate Orbiter and the Mars Polar Lander. Both missions were to satisfy tight financial constraints by exploiting innovative technology under NASA's *faster, better, cheaper* management initiative [571].

The Mars Climate Orbiter was launched in December 1998. It was intended to be the first interplanetary weather satellite. It also had a secondary role to act as a communications relay for the Mars Polar Lander. The Climate Orbiter was to have fired its main engine to achieve an elliptical orbit around Mars in September 1999 [571]. The intention was that it should spend several weeks 'skimming-through' the upper atmosphere. This aero-braking techniques was to achieve a low circular orbit using friction against the spacecraft's solar array to reduce the orbital period from fourteen to two hours. It was during the Mars Orbit Insertion (MOI) maneuver that the Climate Orbiter was lost. The investigation team describe how:

"During the 9-month journey from Earth to Mars, propulsion maneuvers were periodically performed to remove angular momentum buildup in the on-board reaction wheels (flywheels). These Angular Momentum Desaturation (AMD) events occurred 10-14 times more often than was expected by the operations navigation team. This was because the MCO solar array was asymmetrical relative to the spacecraft body as compared to Mars Global Surveyor (MGS) which had symmetrical solar arrays. This asymmetric effect significantly increased the Sun-induced (solar pressure-induced) momentum buildup on the spacecraft. The increased AMD events coupled with the fact that the angular momentum (impulse) data was in English, rather than metric, units, resulted in small errors being introduced in the trajectory estimate over the course of the 9-month journey. At the time of Mars insertion, the spacecraft trajectory was approximately 170 kilometers lower than planned. As a result, MCO either was destroyed in the atmosphere or re-entered heliocentric space after leaving Mars atmosphere. "[565]

The subsequent inquiry identified twelve recommendations for the development and operation of the Polar Lander. These were addressed by the creation of a Mission Safety and Success Team that drew upon fifty of the Jet Propulsion Laboratory's senior staff. A 'red team' was also created to chart all activities that were intended to feed the lessons of the Climate Orbiter incident into the Polar Lander project.

The Mars Polar Lander was launched approximately three months after the loss of the Climate Orbiter in January, 1999. The same cruise stage was to carry the Polar Lander and two smaller probes that were known as Deep Space 2. This was a highly innovative mission that intended to show that miniaturised components could conduct scientific experiments in space. Deep Space 2 consisted of two micro-probes that were to be released from the Polar Lander before it entered the Mars upper atmosphere. These contained a micro-telecommunications system that was designed to communicate with the orbiting Mars Global Surveyor after the probes had impacted with the planet surface. The Polar Lander and the Deep Space 2 probes approached Mars in December 1999. A final trajectory-correction maneuver, TCM-5, was executed six and a half hours before estimated entry. At 12:02 PST, the spacecraft assumed the its entry attitude. A development decision had previously determined that telemetry data would not be collected during the entry, descent and landing phase. In consequence, the change in attitude had the effect of pointing the antenna away from Earth and the signal was lost, as expected. The Polar Lander was expected to touchdown at 00:14 PST and data transmission was scheduled to begin twenty-four minutes later. Data from the DS2 probes was expected to begin at 07:25 No communications were received from either the Polar Lander or the Deep Space 2 probes. The investigation team reported that:

"Given the total absence of telemetry data and no response to any of the attempted recovery actions, it was not expected that a probable cause, or causes, of failure could be determined. In fact, the probable cause of the loss of MPL has been traced to premature shutdown of the descent engines, resulting from a vulnerability of the software to transient

signals. Owing to the lack of data, other potential failure modes cannot positively be ruled out. Nonetheless, the Board judges there to be little doubt about the probable cause of loss of the mission.” [580]

These ‘failure’ of these two missions provides the case study for the remainder of this chapter. A number of motivating factors help to justify this decision. For instance, these incidents provide a rare insight of the way in which organisations must quickly respond to previous incidents. The Jet Propulsion Laboratory and Lockheed Martin had very limited amounts of time to respond to the loss of the Climate Orbiter before the Polar Lander had to be launched. These examples have, however, been deliberately selected for a number of other reasons. They illustrate the failure of leading-edge technology. Previous chapters have shown that the failure of apparently simple technology can be caused by many complex factors. The Allentown explosion discussed in Chapter 8.3 provides an instance of this. The gas line did not rely upon particularly complex technology. However, the incident involved regulatory and organisational failure in the decision not to deploy protective devices and warning systems. The explosion also illustrated complex communication problems between the utility supplier, the excavators, the property owners etc. The immediate causes also reflect a failure in communication and training involving the excavation team and the fire inspectors. The complexity of the modelling in the previous chapter reinforces this meta-level point that even simple technology typically has complex failure modes. In contrast, the loss of the Mars missions provides a completely different challenge. These systems were deliberately designed to ‘push the technological boundaries’ under NASA’s *faster, better, cheaper* management initiative [571].

It is important to address a number of objections that can be made to the inclusion of these incidents. Neither of the Mars Surveyor’98 missions resulted in ‘near misses’. Both involved significant losses in terms of financial resources and in terms of the opportunity costs associated with their scientific objectives. It is important to emphasise, however, that the principle objective in this chapter is to provide readers with a comparative assessment of different analysis techniques. The focus is, therefore, on the analytical techniques rather than the incidents themselves. The same motivations justified the use of the Allentown explosion to illustrate alternative modelling notations in Chapter 8.3. The decision to focus on the Mars Climate Orbiter and the Polar Lander is also justified by NASA’s publication policy. Readers can access a mass of primary and secondary material. I do not know of any near-miss incident that might provide similar opportunities.

Further objections arise because neither of the Mars Surveyor’98 missions posed a direct threat to human safety once it had left the earth’s orbit. It can, therefore, be argued that neither incident is ‘safety-critical’. These two case studies can, however illustrate the application of safety-critical techniques to analyse mission-critical failures. The Mars Climate Orbiter and Polar Lander also illustrate how safety-critical techniques can be applied more generally to understand the causes of technological failure. This is not simply a spurious argument about the theoretical value of safety-critical techniques for mission critical applications. It is a pragmatic observation that has been recognised by many industries. The investigation boards that investigated the loss of the Mars missions were governed by the same regulations that cover investigations into the injury and death of civil-service employees and the general public. NASA Procedures and Guidelines document NPG:8621.1 introduced the term ‘mishap’ to cover these two aspects of mission critical and safety-critical failure [572].

Mission-critical failures provide insights into the possible causes of future safety-critical incidents. This can be seen as a corollary of the previous point. Many analysis techniques reveal common causes of managerial and regulatory failure. As a result, safety and mission-critical incidents may only be distinguished by their consequences rather than by their causes. Leveson reflects this ambiguity when she defines safety to be ‘freedom from loss’ rather than ‘freedom from injury’ [486]. The practical consequences of this have again been recognised by many organisations. For instance, one of the principle findings of the Presidential Commission into the the loss of the space shuttle Challenger was that NASA should establish an Office of Safety, Reliability and Quality Assurance [714]. This agency is intended to have direct authority for safety, reliability, and quality assurance throughout the agency and is independent of other NASA program responsibilities. Such initiatives illustrate the perceived importance of integrating safety concerns into wider quality assurance techniques.

There is little published information about the common causes of safety-related and mission-critical incidents. Previous chapters have mentioned Wright’s preliminary studies, which suggest that accidents may have different causes than incidents [875]. By extension, it can be argued that safety-related incidents may have different underlying causes than mission-critical failures. In particular, it can be argued that mission critical incidents stem from other aspects of dependability, such as security or availability, that have little to do with safety-related failures. Sadly, more time has been spent on debating the semantics of terms such as ‘dependability’ than has been spent on determining underlying differences between mission-critical and safety-critical failure. Much of the discussion focuses on the problems of measuring improvements in such as abstract notion when it can be influenced by many more detailed factors including reliability, safety, security, availability etc [475, 486]. For example, a security improvement might increase the dependability of a system in some abstract sense. It can also jeopardise safety if operators are prevented from accessing necessary functions during a systems failure. This debate reflects divisions within the academic community. It also reflects pragmatic distinctions that shape organisational responses to technological failure. For example, NASA’s Office of Safety and Mission Assurance provides a common focus for dependability concerns. This organisation does not, however, derive abstract measures of dependability. The focus is on gathering and analysing more detailed information about the causes of mission success and failure. Brevity prevents a more detailed analysis of the practical implications of distinctions between the various components of dependability. In contrast, our focus is on determining whether similar analytical techniques can provide insights into both safety-critical and mission-critical incidents. At present there is insufficient evidence to prove or disprove this hypothesis. The case studies in this chapter can, however, be usefully compared to previous work in incident analysis [411, 470]. Although the analysis presents a single view upon two isolated case studies, there are many strong similarities between the detailed causes of these mission failures and the causes of safety related incidents that were identified in Chapter 2.3. This should not be surprising given that these safety-related factors are often presented as generic causes of technological and managerial failure.

10.2 Stage 1: Incident Modelling (Revisited)

This section introduces what the US Department of Energy has described as the ‘core’ analytical techniques for incident and accident investigation. Figure 10.1 provides an overview of these techniques.




Figure 10.1: Overview of the Dept. of Energy’s ‘Core’ Techniques

The following pages focus on the modelling techniques that form a precursor to any subsequent causal analysis. In order to understand *why* an incident occurred, it is first necessary to determine *what* happened. These are illustrated on the left hand side of Figure 10.1. Unfortunately, the expressive power of these modelling notation is not as great as some of those introduced in Chapter 8.3. As we shall see, it can be difficult to represent and reason about detailed temporal relationships

between the events that are represented in these ‘core’ modelling techniques. With these caveats in mind, the following sections show how event and causal analysis charts can be used to represent the products of barrier and change analysis. The resulting diagrams then support a more detailed root cause analysis.

10.2.1 Events and Causal Factor Charting

Event and Causal Factor (ECF) charts provide a graphical means of representing the sequence of events leading to a failure. These charts are then annotated with additional causal information. For now, however, it is sufficient to observe that the motivating factors that justify the maintenance of these charts are the same as those for the techniques introduced in Chapter 8.3:

“Constructing the events and causal factors chart should begin immediately. However, the initial chart will be only a skeleton of the final product. Many events and conditions will be discovered in a short amount of time, and therefore, the chart should be updated almost daily throughout the investigative data collection phase. Keeping the chart up to date helps ensure that the investigation proceeds smoothly, that gaps in information are identified, and that the investigators have a clear representation of accident chronology for use in evidence collection and witness interviewing.” [208]

Figure 10.2 provides a high-level view of the components of an events and causal factor chart. A number of guidelines support the development of these diagrams [208]. The process begins by mapping out a chronology of events. Time is assumed to flow from the left of the diagram to the right. Events represent actions and should be stated with one noun and one active verb. They should be quantified “as much as possible and whenever applicable”. The examples suggest that analysts specify how far a worker falls rather than only state that the fall occurred. Times and dates must also be noted and the events should “be derived from” the events that precede them. The approach, therefore, has strong similarities with the use of timelines in previous chapters. Analysts must, however, also distinguish a primary chain from other sequences of events that contribute to the failure. These secondary chains are drawn above the primary line. Without tool support, the problems of maintaining complex graphical structures can limit the scope for introducing these additional event sequences.




Figure 10.2: Simplified Structure of an ECF Chart

As mentioned, ECF charts have a superficial similarity to timelines. Both exploit linear structures to denote the flow of events leading to an incident or accident. Both approaches must, therefore, consider how to represent state-based information and emergent properties that develop slowly over time. In the case of ECF Charts, these are denoted by the conditions that appear in the ellipses of




Figure 10.3: Components of ECF Chart

Figure 10.2. Conditions are passive. For example, they denote that ‘there was bad weather’ or that ‘workers were tired’. They are also associated with the particular events that they help to influence.

Figure 10.3 presents the component symbols that are used in ECF Charts. As with our use of modelling notations, this approach needs to be adapted to support incident analysis. For instance, the diamond used to denote an accident in Figure 10.3 can be used more generally to represent the potential outcome of a ‘near miss’ incident. Similarly, it is likely that there will be far more presumptive events and conditions in certain types of incident report systems. For example, analysts are more likely to be forced to make inferences about the events leading to an incident if they have to piece together information from a single submission to an anonymous system. Figure 10.4 illustrates how the ECF notation can be applied to represent the loss of the Mars Climate Orbiter. The intention is to illustrate the information that might be available to investigators in the immediate aftermath of an incident. As can be seen, the primary flow of events is assumed to begin with the launch of the mission on the 11th December. Subsequent analysis will extend the scope of events to consider decisions that were made prior to launch. However, such information may not immediately be available immediately after such an incident. The mission progressed until the last signal was received at 09:04:52.




Figure 10.4: High-Level ECF Chart for the Mars Climate Orbiter (MCO)

A number of comments can be made about the use of the ECF notation in Figure 10.4. The accident symbol is used to denote the loss of the Climate Orbiter; MCO is lost. It does not describe the nature of the incident in great detail. NASA investigators considered two possible scenarios; either the craft was destroyed in Mars' atmosphere or it re-entered heliocentric space. These are not shown here because we do not know whether these possible incidents actually took place. This ambiguity stems from NASA's decision not to relay telemetry data during Mars Orbit Insertion. The same decision was taken during the development of the Polar Lander. This deliberate design feature reduced project development costs but clearly also reduced the information that was available to subsequent investigators. As the analysts commented "the decision not to have EDL telemetry was a defensible project decision, but an indefensible programmatic one." [580].

A second important feature of Figure 10.4 is the way in which it extends beyond the loss of the MCO's signal. The Operational Navigation team met with Spacecraft Engineers to discuss what might have caused the apparent mission failure. This meeting formed part of an initial response that was intended to devise a way of re-establishing contact with the mission and then, later, to learn any immediate lessons that might affect the Mars Polar Lander. Shortly after this meeting, a bug was discovered in the 'Small Forces' software that formed an important component of the navigation system. This sequence of events is critical to any understanding of the MCO incident, not simply because it helped to identify the probable cause of the failure but also because it took place *before* the NASA Mishap investigation board had been formed.

It is inevitable that informal analysis will be conducted in the aftermath of many incidents. In particular, the limited launch window for the Mars Polar Lander made it imperative that lessons were learned as quickly as possible. It can also be argued that by discussing the causes of failure, engineers can make the best use of any opportunities to mitigate the consequences of an incident. However, there also a number of concerns about such interim forms of analysis. Firstly, operators may actually exacerbate the situation if they intervene with partial knowledge about the causes of an incident. The Chernobyl and Three Mile Island accidents provide graphic illustrations of this point. In the former case, Soviet operators exacerbated their problems by rapidly inserting control rods into the reactor that had previously been almost fully withdrawn. Rather than dampening the reaction, positive void coefficients created the opposite effect. Operator intervention at Three Mile Island led the NRC to specify that users should not intervene in similar circumstances without a sufficient period to formulate a detailed diagnosis of the causes of the failure [220]. Secondly, there is a danger that groups who are involved in an incident may prepare an explanation of the failure that cannot be supported by a more detailed analysis. At its most extreme, this may extend to collusion in falsifying evidence. At its most benign, the identification of a probable cause by groups of workers in the aftermath of an incident can have the effect of biasing, or blinkering, any subsequent investigation. Neither of these objections can be applied to the MCO engineers or to NASA's Mishap Investigation board. It should be noted, however, that the MCO phase I report focuses almost exclusively on the faults identified by the Operational Navigators and the Spacecraft Engineers following their meeting on the 27th September.

Figure 10.5 extends the previous ECF chart to illustrate an interim stage in the analysis of the MCO incident. As can be seen, this diagram focuses in on events between the launch and the completion of the cruise phase. In particular, it focuses on Angular Momentum Desaturation events. These maneuvers were partially determined by the 'Small Forces' software. As Figure 10.4 shows, this was the code that had been identified as the potential problem by the Operational Navigators and the Spacecraft Engineers. Figure 10.5 shows that ground based software used pounds of force per second rather than Newtons per second to represent thruster performance. This code was used to generate the Angular Momentum Desaturation file that was then used as input to subsequent navigation software and so repeated AMD events would compound any inaccuracies. The condition above the AMD event denotes the observation that Angular Momentum Desaturation maneuvers had to be carried 10 to 14 times more often than had been planned. This was to counter-act the momentum that was induced by radiation acting on the spacecraft's solar array. As can be seen, a secondary line of events explains why AMD maneuvers were so common. A decision was taken to use asymmetric solar panels. this was different to the symmetric configuration used on the Mars Global Surveyor. The frequency of AMD events on the MCO also stemmed from a decision not to




Figure 10.5: Angular Momentum Desaturation Events Affect MCO Navigation

perform what were termed ‘barbecue’ maneuvers in which the craft was flipped through 180 degrees every twenty-four hours.




Figure 10.6: High-Level ECF chart for the Mars Polar Lander (MPL)

Previous ECF charts have focussed on the loss of the MCO. In contrast, Figure 10.6 presents a very high-level view of the observable events that took place before the loss of the Mars Polar Lander. It is important to note again that this diagram does not represent the exact events that might have contributed to the loss of the Lander and the Deep Space 2 probes. The Mars Polar Lander and Deep Space 2 missions might have been destroyed in the atmosphere or re-entered heliocentric space. They might also have been damaged by impact on landing or communications failures might have prevented subsequent communication. The lack of telemetry data can prevent analysts from assessing the likelihood of these different scenarios until a secondary investigation is completed. It is also important to note that this incident is slightly more complex than the loss of the Climate Orbiter. Any failure scenario represented by an ECF chart must account for the loss of the Lander as well as both of the Deep Space 2 mission. Both probes could independently communicate with the Mars Global Surveyor after they had been deployed on the planet surface. A single failure mode is most likely to have occurred prior to the separation of the probes from the Lander. Any failure after separation is most likely to have involved two different failure modes.

Figure 10.7 provides a more detailed view of two of the failure modes that might explain the loss of the Polar Lander and Deep Space 2 missions. As can be seen, the nature and scope of the ECF chart will change as more information becomes available. In this example, the loss of the Polar




Figure 10.7: Premature MPL Engine Shut-Down and DS2 Battery Failure

Lander occurs after the premature shut down of the engines at forty feet from the planet surface. This is influenced by a software condition which specified that the engines should be cut if there were two consecutive readings from Hall effect magnetic sensors and the Lander's radar detected that the surface was less than forty meters away. Hall effect sensors were attached to each of the Lander's legs. These were intended to function as follows. Once a leg touched the surface of the planet, the resultant motion would move a magnet away from the sensor. This movement would reduce the magnetic field below the sensor's trigger level. However, as can be seen from the upper-left event in Figure 10.7, spurious signals are generated by the sensors when the legs are first deployed into a landing position at some 1,500 meters from the surface. To prevent this from having a disastrous effect, the software systems disregard any signals that are received from the Hall effect sensors until the on-board radar detects the surface at less than forty meters above the surface. The ECF chart in Figure 10.7 represents a possible failure sequence for this approach. If the sensors generate two consecutive spurious signals on leg deployment then a variable Touchdown is initially marked as true. This is not reset to False even though the on-board radar detects that the surface is more than 40 meters away. As a result, when the radar eventually does detect that the surface is 40 meters away the software retains the spurious value of the Touchdown signal that was generated during leg deployment. The two conditions in the software are now satisfied and the engines are cut even though none of the legs are in contact with the surface.

Figure 10.7 also represents different events leading to the loss of the Deep Space 2 probes. These probes would have separated from the Lander long before the engines were cut and so a different explanation has to be found for the loss of any signal between these devices and the Mars Global Surveyor. A presumptive event is used to denote that the probes correctly separated from the Lander. There is no means of being completely sure that this did occur given the lack of telemetry data. A number of alternative failure scenarios can be considered in which the separation did not take place, these would have to be represented in additional ECF chart. In this example, however, correct separation leads to the assumptions that the probes impacted with the planet surface but that both suffered an electrical failure. The associated condition is used to indicate that this is a possible failure scenario because there are no common mode failures in the penetrator section of the probe that could cause a failure in the telecommunications systems. This is a slight simplification if the tethering mechanisms is considered to be part of the penetrator. The loss of both probes can be explained by a failure in either the radio assembly or the battery components that were both located in their aft section.

It is important to stress that the ECF charts in this section provide a very limited view of the possible failure scenarios. In practice, investigators must develop a number of similar diagrams to represent alternative sequences of events. It is important also to remember that the ECF technique was not initially intended to support the analysis of high-technology failures within the aerospace industry. The Polar Lander and Climate Orbiter case studies were deliberately chosen as a challenge to the application of these analytical techniques. For example, the decision not to provide telemetry links during the Lander's Entry, Descent and Landing or the Orbiter's insertion creates a degree of uncertainty that is not often apparent in the more usual application of ECF diagrams to occupational injuries [208].

This section has shown how ECF charts can be used to develop high-level reconstructions of the events that contribute to particular failure scenarios. As can be seen, this involves the identification of observable events, such as the last signals from the Lander, and presumptive events, such as battery damage to the Deep Space 2 probes. These diagrams, therefore, represent an initial stage in the causal analysis of an incident [210]. However, they do not go much beyond the reconstructive modelling techniques that were introduced in Chapter 8.3. To distinguish between root causes and contributory causes, investigators must recruit a range of complementary analytical techniques. These can be used to ask deeper questions about *why* particular events did or did not contribute to a failure scenario. The results of techniques, such as barrier analysis, can then be used to develop more detailed ECF diagrams.

10.2.2 Barrier Analysis

Barrier analysis has its modern roots in the early 1970's when Haddon proposed a taxonomy of different controls that can be used to mitigate or direct the transfer of energy in safety-critical systems [300]. These included measures to reduce the amount of energy that is generated, measures to separate a target from the source of energy either in time or space, measures to modify shock concentration surfaces and to strengthen the target. These general ideas led to the development of more formal techniques for barrier analysis both as a tool for incident analysis and also as a constructive design tool. As with ECF charting, this technique was driven by the requirements of the US Department of Energy to develop techniques that support the development and analysis of a range of hazardous processes, including nuclear power generation. It is important to stress that barrier analysis also supports the reconstruction and simulation techniques that were described in previous chapters. Fault trees, time-lines, Petri Nets can all be used to capture insights about the successes and failures of potential 'protection devices'. However, barrier analysis is most often used by analysts as a means of extending an initial ECF chart to consider a broader range of potential root causes.

Barrier analysis starts from the assumption that a hazard comes into contact with a target because barriers or controls were unused or inadequate. A hazard is usually thought of as an unwanted energy transfer such as the passage of electricity from an item of equipment to an unprotected worker. Energy can be 'kinetic, biological acoustical, chemical, electrical, mechanical potential, electro-magnetic, thermal or radiation' [208]. The target is the person, equipment or other object that can be harmed by a hazard. Barriers represent the diverse physical and organisational measures that are taken to prevent a target from being affected by a potential hazard. Although distinctions are blurred, many barrier analysis techniques identify controls and safety devices. Control barriers direct wanted or 'desired' energy flows. They include conductors, disconnect switches, pressure vessels and approved work methods. Safety devices are barriers to unwanted energy flows. These include protective equipment, guard rails, safety training and emergency places [209]. The reason that such distinctions can be difficult to make is that the same energy flow might be both wanted and unwanted at different times during an application process. For instance, the Landers thrusters deliver necessary power during the landing sequence. However, this same power source might topple the craft if it continues after the legs have touched the planet surface. The Hall sensors can, therefore, be seen both as controls and safety devices. They acted as a control during the descent because they kept the thrusters working. If the engines were cut then the Lander would be destroyed. However, once the craft has landed the same devices act as safety devices because the power is no longer wanted. Have acknowledged the practical difficulties created by any distinction between safety and control devices, it is possible to distinguish a number of further barriers.

It is possible to identify three different forms of barriers: people; process and technology. For example, material technology has produced physical barriers that directly prevent a hazard from affecting a target. They include guards, gloves and goggles, protective clothing, shields. As we shall see, these devices are often rated to be effective within certain tolerances. For example, a fireguard may provide protection against a fire within particular heat and time limitations. Dynamic barriers include warning devices and alarms [209]. These are not continually apparent but are only issued when the system detects that there may be a potential hazard. This definition can also be extended to include physical interlocks that restrict access or actions during critical phases of an operation. The limitations with this approach stem from the dynamic nature of these warnings. Operators may fail to notice information about a potential hazard. Operators may also choose to disregard or circumvent warnings, especially, if they have been presented with a succession of false alarms. Conversely, warnings may not be invoked even though a hazard may be present. This poses a particular threat if operators grow accustomed to the additional protection afforded by these barriers.

Process barriers include the use of training, of checklists, of standard operating procedures and other forms of workplace regulation that are intended to protect operators and their equipment from potential hazards. Chapter 2.3 has argued that these procedures can either be explicitly supported by line management or they may arise over time as the result of implicit procedures within everyday

working practices. The later class of barriers can be unreliable if new employees fail to observe the way in which existing employees follow these unwritten rules.

People also represent a further class of barrier that can protect a target from a hazard. Human often act as the last barrier against the adverse consequences of energy transfers. The Office of Operating Experience, Analysis and Feedback in the US Department of Energy concludes that:

“Human action is often, but not always, associated with a procedural barrier. Examples of human action serving to control a hazard are controlling and extinguishing a fire, de-energizing an electrical circuit either in response to a procedure or as part of safe work practice, evacuating a building in response to a fire or a criticality alarm, etc.” [206].

Managerial and administrative policies can also be interpreted as a form of meta-level barrier. These constraints do not directly protect any particular target from any particular hazard. For instance, they do not directly involve a physical device shielding the operator from a heat source. In contrast, managerial and administrative barriers help to ensure that the acquisition, development, installation and maintenance of a system ensures the adequate provision of more direct barriers to protect potential targets.

The previous paragraphs have mentioned that there are a number of different ways in which barriers can fail. The following list provides a high-level overview of these failure modes:

- *Barrier is impractical - impossible.* There are situations in which it is impossible to provide adequate barriers against a potential energy transfer. Ideally, such situations are identified during a safety analysis. If the hazard could not be prevented or mitigated, regulators should ensure that the process fails to gain necessary permissions. Payne provides numerous examples of this in his analysis of planning applications for safety-critical production processes [671]. He cites a series of incidents in which it was impossible to protect the public once chemicals had been released into the environment. In retrospect, permission should not have been granted for the processes to be sited within urban developments.
- *Barrier is impractical - uneconomic.* In other circumstances, it may be technically feasible to develop appropriate barriers but their cost may prevent them from being deployed. As we have seen, a spate of ‘near misses’ and accidents persuaded regulators to back the introduction of a Train Protection Warning System on UK railways. This is estimated to cost approximately £310 million. The more sophisticated Advanced Train Protection system was rejected as being uneconomic, at an estimated cost of £2 billion [691]. The obvious weakness with this form of analysis is that the perceived benefits that are associated with particular barriers can change in response to public anxiety over particular incidents. The Southall and Paddington crashes led to a detailed reassessment of the economic arguments against the introduction of the more advanced system.
- *Barrier fails - partially.* A barrier that has been successfully introduced into an application process may, however, fail to fully protect the target from a potential hazard. This is an important class of failure in many incident reporting systems because it represents situations in which barriers provide some protection but may not, under other circumstances, have prevented the hazard from being released. For instance, the Mishap Investigation Board into the loss of the Climate Orbiter directed the Polar Lander team to introduce a series of protective barriers. These included the establishment of a ‘red team’ that was intended to:

“study mission scenarios, to ensure operational readiness and to validate risks... This team provides an independent, aggressive, almost adversarial yet helpful role, addressing all levels of the project from high-level requirements down through sub-system design. Key review items include: ensuring system success and reliability; reviewing overall system design and design decisions; reviewing system safety and reliability analyses and risk assessments; reviewing planned and completed testing; and reviewing operational processes, procedures and team preparation. Red team

review results and recommendations are reported to the project manager and the project team, as well as senior level management at the centers.” [571]

While this device undoubtedly helped to protect the Polar Lander against a number of potential hazards, it failed to provide total protection against the failure modes that were identified in the aftermath of this second incident.

- *Barrier fails - totally.* The distinction between partial and total protection depends upon the nature of the application. This can be illustrated by assuming for a moment that the failure scenario in Figure 10.7 is an accurate representation of the events leading to the loss of the Polar Lander. The on-board systems prevented it from immediately cutting its engines when the Hall effect sensors first detected spurious readings. From this perspective, the software provided partial protection. However, the software completely failed in terms of the overall mission objectives. The protection was insufficient to ensure the safe landing of the craft. This example illustrates how the success or failure of a barrier must be interpreted with respect to the overall safety objectives of the system as a whole. The craft was lost and hence the protection is interpreted to have failed in its intended function.
- *Barrier is not used - not provided.* This describes a situation in which a barrier might have protected a target had it been available. At a prosaic level, the bug in the Polar Lander software could have been removed by the addition of a statement, (`IndicatorState = False`), when the radar detects the forty meter threshold. This need not have provided total protection for the mission. There are a number of alternative failure modes. For instance, the Lander may have encountered terrain with a slope steep enough to destabilize the craft on landing.
- *Barrier is not used - by error.* Barriers may not be used during an incident even though they are available and might prevent a target from being exposed to a hazard. For example, the Climate Orbiter had a contingency maneuver plan in place to execute a Trajectory Correction Maneuver (TCM5). This was intended to raise the the orbit, in fact the second periapsis passage, to a safe altitude [571]. TCM5 could have been used shortly before Mars Orbit Insertion as an emergency maneuver. It was discussed verbally before the MOI but was never executed. The NASA investigators commented that “the analysis, tests and procedures to commit to a TCM5 in the event of a safety issue were not completed, nor attempted” [571]. In consequence, the operations team were not prepared for such a maneuver.

The previous paragraphs have introduced a number of high-level concepts: barriers; targets and hazards. We have also identified ways in barriers may fail to protect a target or may not be available to mitigate or control a potential hazard. We have not, however, provided a mechanism by which these general observations can support the causal analysis of adverse occurrences. Nor have we shown how the findings of such an analysis can be integrated into the ECF charts that were developed in the previous section. Barrier tables, such as that shown in Table 10.1, can be used to address this omission.

Table 10.1 provides a high level view of the barriers that were intended to prevent the Climate Orbiter from re-entering heliocentric space or impacting the planet surface. As can be seen, the people, process and technology distinctions are retained from the previous paragraphs. This reflects the key components for *Mission Success First* that was advocated by the NASA mishap investigators. They argued that “every individual on the program/project team (must) continuously employ solid engineering and scientific discipline, take personal ownership for their project development efforts and continuously manage risk in order to design, develop and deliver robust systems capable of supporting all mission scenarios” [571]. Table 10.1 records some of the reasons why the individuals involved in the Climate Orbiter project failed to adequately protect against the potential loss of the mission.

People Barriers

Firstly, there were insufficient staff. The primary investigation found that the staffing of the operations navigation team was less than adequate. In particular, the Mars Surveyor Operations Project

Hazard: Impact/Re-Entry	Target: Mars Climate Orbiter
Barrier	Reason for failure?
People	Lack of staff
	Changes in management
	Inadequate training/skills
	Poor communication
Process	Separation of development and operations teams
	No systematic hazard analysis
	Inadequate testing
	Lack of oversight
Technology	Incorrect trajectory modelling
	Tracking problems
	Rejection of barbecue mode
	Rejection of TCM-5

Table 10.1: Level 1 Barrier Table for the Loss of the Climate Orbiter.

was responsible for running the Global Surveyor and the Polar Lander in addition to the Climate Orbiter. The investigation revealed that these divided responsibilities tended to ‘dilute’ the focus on any single mission. This loading had a particular effect on the Climate Orbiter’s navigation team. The two individuals who led this group found it very difficult to provide the twenty-four hour a day coverage that was recommended during critical phases of a mission, such as the Climate Orbiter’s MOI [565]. The loss of the Climate Orbiter led to an increase in the number of navigators who were assigned to the Polar Lander project. In terms of the earlier mission, however, this lack of personnel may have prevented the navigation team from sustaining their investigation into the anomalies that they found between the ground-based and on-board navigation systems. This, in turn, reduced the navigation team’s ability to operate as an effective barrier to any navigational problems that might ultimately threaten the success of the mission.

Barrier analysis can also be used to identify further ways in which individuals failed to prevent the loss of the Climate Orbiter. In particular, changes in management prevented an effective response to the navigation problems. During the months leading up to MOI, the investigators found that the Mars Surveyor operations team had “some key personnel vacancies and a change in top management” [571]. A number of further problems reduced management effectiveness in combating particular hazards. For example, there was a perceived ‘lack of ownership’ by some operations personnel who felt that the mission had simply been passed onto them by the development teams. A key management failure in this process was that the operations team had no systems engineering or mission assurance personnel who might have monitored the implementation of the process. This, in turn, might have helped to improve communication between these different phases of the mission. Poor communication appears as a separate explanation for the way in which human barriers failed to prevent mission failure. The investigators concluded that “the spacecraft operations team did not understand the concerns of the operations navigation team” [565]. The operations navigation team appeared to be isolated from the development team and from their colleagues in other areas of operations. Other problems stemmed from the nature of group communications during the cruise phase. For example, the navigation team relied on email to coordinate their response once the conflicts were identified in the navigation data. The investigators were concerned that this use of technology enabled some of the problems to ‘slip through the cracks’.

Primary and secondary investigations also identified inadequate training as a potential reason why staff failed to identify the potential hazard to the mission. This was connected to the lack of key personnel because there was no adequate means of ensuring that new team members acquired necessary operational skills. In particular, there was no explicit mentoring system [571]. The investigators argued that the “failure to use metric units in the coding of the Small Forces ground

software used in trajectory modeling...might have been uncovered with proper training” [565]. Such comments are significant because they come very close to the counterfactual arguments that have been associated with root cause analysis [25]. One particularly important area for concern was that the operations navigation team was not familiar with the attitude control system on-board the Climate orbiter; “these functions and their ramifications for Mars Climate Orbiter navigation were fully understood by neither the operations navigation team nor the spacecraft team, due to inexperience and miscommunication” [571]. This lack of familiarity with spacecraft characteristics had considerable consequences throughout the incident. In particular, it may have prevented the operational navigation team from appreciating the full significance of the discrepancies that were identified.

Table 10.1 summarises the reasons why individuals failed to protect the Climate Orbiter from mission failure. The previous paragraphs have built upon this analysis to explain why lack of staff, changes in management, inadequate training and poor communication had an adverse effect upon potential barriers. We have not shown how the results of this analysis might be used to inform the development of Effects and Causal Factor diagrams. The first problem in incorporating these additional insights is that many of the barriers, described above, relate to distal factors. They influence several of the events in Figures 10.4 and 10.5. A second issue is that barrier analysis, typically, helps to identify additional events that ought to be introduced into an Effects and Causal Factor diagram. This is particularly important because primary investigations often focus on catalytic events rather than events that weakened particular barriers.




Figure 10.8: Integrating the Products of Barrier Analysis into ECF Charts

Figure 10.8 integrates our analysis of the human barriers to mission failure into an ECF chart. As can be seen, this diagram introduces a new event into the primary sequence. This denotes the decision not to initiate the TCM-5 maneuver. It was introduced because the previous barrier analysis identified TCM-5 as an important opportunity for preventing the hazard from affecting the target. Figure 10.8 also uses the insights from the barrier analysis to explain why this opportunity was not acted upon. Lack of staff, inadequate training, management changes and poor communication between the operational navigation and spacecraft teams were all factors in the failure to perceive the significance of the AMD data anomaly. Figure 10.8 also illustrates the way in which barrier analysis helps to identify key event sequences that may not have been identified during the initial analysis of an adverse occurrence. As can be seen, this ECF chart has been extended to represent the fact that file formatting errors prevented the navigation team from identifying the AMD anomaly until more than four months after launch.

Process Barriers

Table 10.1 identified four ways in which process barriers may have failed during the Climate Orbiter incident. These related to the separation of the development and operations teams, to the lack of any systematic hazard analysis, to inadequate testing and to the lack of management oversight during particular phases of the mission.

The previous section identified that many of the operational staff lacked necessary training about the operating characteristics of the Climate Orbiter. One reason for this was that the overall project plan did not provide for a careful hand-over from the development project to the operations staff. The Climate Orbiter was also the first mission to be supported by a multi-mission Mars Surveyor Operations Project. The operations staff had to assume control of the Climate Orbiter project without losing track of the Global Orbiter and the Polar Lander missions. These logistical problems were compounded by that fact that the Climate Orbiter project was the first Jet Propulsion Laboratory mission in which only a small number of development staff were ‘transitioned’ into the operations team. No navigation personnel, made this move from the development of the Climate Orbiter into its operation. This had a number of important consequences for subsequent events during the incident. In particular, the navigation team and other operational staff may have made a number of incorrect assumptions about hardware and software similarities between the Global Surveyor and the Climate Orbiter. The investigators argued that:

“This apparently caused the operations navigation team to acquire insufficient technical knowledge of the spacecraft, its operation, and its potential impact to navigation computations. The operations navigation team did not know until long after launch that the spacecraft routinely calculated, and transmitted to Earth, velocity change data for the angular momentum desaturation events. An early comparison of these spacecraft-generated data with the tracking data might have uncovered the units problem that ultimately led to the loss of the spacecraft.” [565].

The key point here is that the decision not to transition key development staff into the operation phase removed one of the procedural barriers that otherwise protect JPL missions. The navigational operations team might have realised the potential significance of the AMS anomaly if they had known more about the decisions that had informed the development of the Climate Orbiter.

Figure 10.9 shows how barrier analysis helps to identify a number of additional events and conditions that influenced the course of the incident. The ECF chart has been extended to explicitly denote that a minimal number of development staff were transferred to the operations teams. A number of associated conditions show that the plans for this transition were less than adequate and that this was the first project for the multi-mission Mars Survey Operations project. The previous barrier analysis, however, also raises a number of important questions about the construction of ECF charts. For example, the decision only to transfer a minimal number of staff helped to create the conditions in which operational teams made inappropriate assumptions about the similarity between the Global Surveyor and the Climate Orbiter. These erroneous nature of these suppositions is underlined by the changes in the solar array that are also noted on Figure 10.9. Problems arise




Figure 10.9: Process Barriers Fail to Protect the Climate Orbiter

because although these incorrect assumptions stem from early in the transition from development to operations, they continue to have an influence throughout the incident. This is difficult to denote use the ECF format introduced in previous section. The condition that represents the potential for incorrect assumptions is surrounded by a double line. Later sections will explain how such conditions provide an important starting point for any subsequent attempts to distinguish root causes from contributory factors.

The hand-over from development to operation was one of several process issues that undermined the Climate Orbiter mission. The lack of any systematic hazard assessment, for instance using Fault Tree analysis, had numerous consequences for the mission as a whole. This prevented engineers from considering a range of possible failure modes. It also prevented the development and operations teams from conducting a systematic assessment of what were, and what were not, mission critical features. In particular, some form of hazard analysis might have helped to identify that specific elements of the ground software could be ‘mission critical’ for the operations navigation team. Finally, the lack of a coherent hazard analysis may also have led to inadequate contingency planning. This is particularly apparent in the lack of preparation for TCM-5, mentioned in previous paragraphs. As can be seen, the failure to conduct such an analysis had the knock-on effect of removing a number of potential barriers that might have either detected the navigation software as a critical component prior to launch or might, subsequently, have encouraged operations to reconsider contingency plans once the anomaly had been discovered.

The previous paragraph argued that the lack of any systematic hazard analysis illustrates a




Figure 10.10: Process Barriers Fail to Protect the Climate Orbiter (2)

further failure of process barriers. Figure 10.10 builds on this analysis by integrating it into the previous ECF charts. This illustrates one of the issues that can complicate the construction of such diagrams. It can be difficult to decide whether or not a particular failure should be represented by the event that triggered the failure or by the conditions that form the consequences of that event. For example, Figure 10.10 include an event labelled Decision not to perform an a priori analysis of what could go wrong on the MCO. This might have been represented by a condition labelled there was no systematic hazard analysis. The ECF manuals provide little guidance on this issue [210, 208]. It is important, however, that some heuristic be used to guide the construction of these diagrams. We have, therefore, use events to denote those stages in an incident that might become a focus for subsequent analysis. Investigators might decide that more needs to be known about the circumstances that influenced any decision not to conduct a systemic hazard analysis. This decision is, therefore, represented as an event rather than a condition.

Further process barriers were undermined by the lack of any sustained validation at a systems level. Navigation requirements were set at too high a management level. In consequence, programmers and engineers were left to determine how best to satisfy those requirements without detailed guidance from others involved in the development process. These problems might not have been so severe had their consequences been detected by an adequate validation process. Several significant system and subsystem flaws were, however, only uncovered after the Climate Orbiter had

been launched. For instance, file format errors prevented the navigation team from receiving and interpreting telemetry from the ground system for almost six months. The NASA investigators argued that there was “inadequate independent verification and validation of Mars Climate Orbiter ground software (end-to-end testing to validate the small forces ground software performance and its applicability to the software interface specification did not appear to be accomplished)” [571].

The validation issues and the lack of any system level hazard analysis were exacerbated by a more general lack of oversight during the Climate Orbiter mission. There was little Jet Propulsion Laboratory oversight of Lockheed Martin Astronautics subsystem developments. This created problems as the level of staffing was reduced during the transition from development to operations. Several mission critical functions, including navigation and software validation, received insufficient management oversight. It also became difficult to maintain lines of responsibility and accountability during the project. This point can be illustrated by the Mishap board’s description of the relationship between JPL and the contractor:

“Lockheed Martin Astronautics of Denver, Colorado was selected as the prime contractor. Lockheed Martin Astronautics contracted development responsibilities were to design and develop both spacecraft, lead flight system integration and test, and support launch operations. JPL retained responsibilities for overall project management, spacecraft and instrument development management, project system engineering, mission design, navigation design, mission operation system development, ground data system development, and mission assurance. The Mars Surveyor Project’98 assigned the responsibility for mission operations systems/ground data systems development to the Mars Surveyor Operations Project, Lockheed Martin Astronautics provided support to Mars Surveyor Operations Project for mission operations systems/ground data systems development tasks related to spacecraft test and operations.” [565]

Recurring questions in the NASA investigation included ‘Who is in charge?’ and ‘Who is the mission manager?’. The investigators reported repeated examples of ‘hesitancy and wavering’ whenever individuals attempted to answer the latter question. This is not surprising given the comments made about the feelings of guilt and blame that often operators’ reactions to adverse occurrences, see Chapter 4.3. However, the NASA board also describe how one interviewee answered that the flight operations manager was acting like a mission manager without being designated as such.

Figure 10.11 shows how the insights that can be derived from a barrier analysis of process failures can be represented within the previous ECF charts. As can be seen the lack of oversight had an important effect on many diverse aspects of the Climate Orbiter’s development and operation. If this oversight had been in place then it might have persuaded participants to be more circumspect in their assumptions about the Climate Orbiter’s hardware and software characteristics. More coherent oversight might also have encouraged a systemic hazard analysis, especially if more attention had been paid to the validation of high-level requirements.

It should be apparent from the preceding paragraphs that there is no automatic means of propagating the findings of a barrier analysis into the graphical representations of an ECF chart. The investigator must determine how best to translate the findings of their analysis into the events and conditions of Figures 10.10 and 10.11. It, therefore, follows that different investigators might derive different event structures from those shown in this chapter. This introduces a number of concerns about the consistency and validity of any analysis. I am unaware of any research having been conducted into these important aspects of the ECF technique. It can, however, be argued that this analytical process is less about the development of a single coherent view than it is about the explicit representation of what might otherwise remain implicit assessments about the success or failure of particular barriers.

Technological Barriers

Technological barriers can also be deployed to support the protection that people and processes provide for safety-critical and mission-critical applications. Table 10.1 has identified four ways in which these technological barriers failed to support the Climate Orbiter mission. There were




Figure 10.11: Process Barriers Fail to Protect the Climate Orbiter (3)

problems with the trajectory modelling that was intended to identify that potential navigation hazards. The tracking systems that were intended to identify failures in the trajectory models also provided contradictory information. The failure of these barriers became increasingly important because of decisions not to exploit some of the technological measures, including the barbecue mode and TCM-5 contingency, that might otherwise have prevented the mishap from occurring.

The barbecue mode involved a plan to ‘flip’ the spacecraft by 180 degrees every twenty-four hours. This would have reduced the need for AMD events. The rotation of the aircraft would ensure that any momentum induced by the asymmetric solar panels would have been counteracted in the following twenty-four hours. Previous sections have already shown how this decision can be introduced in an ECF chart, for example Figure 10.5. Similarly, Figure 10.8 introduced the decision not to initiate the TCM-5 maneuver into previous ECF charts. This formed part of an analysis into the failure of people-related barriers. Rather than extend the scope of these previous diagrams, this section focuses on the technological problems that removed navigation and tracking safeguards. Subsequent paragraphs go on to perform a more detailed analysis of the software ‘bugs’ that removed many of the technological barriers to mission failure.

The previous section has described how problems in the validation of mission critical software created a situation in which several systems had to be debugged during the cruise phase of the mission. This created particular problems because these systems provided important barriers against mission failure. In particular, ground software could not be used to perform the anticipated Angular Momentum Desaturation calculations during the first four months of the cruise. Multiple file format errors were compounded by problems with the data types that were used to represent the spacecraft’s attitude. As we have seen, the operations navigation team was forced to use email from the contractor to notify them when a desaturation event was occurring. They then attempted to model the impact on the Climate Orbiter’s trajectory using timing information and the manufacturer’s performance data. It was not until April 1999 that operations staff could begin using the correctly formatted files. It took a further week for the navigation team to diagnose that the files underestimated the

trajectory perturbations due to desaturation events.

The file format and content errors removed important barriers that might otherwise have protected the mission. They prevented the operations navigation team from being able to quickly detect and investigate the underlying calculation problems. These problems might not have had severe consequences if other forms of protection had also been available. In particular, the operations navigation team had limited means of tracking and monitoring the consequences of AMD events. It was difficult to observe the total magnitude of the thrust because of the relative geometry of the thrusters used for AMD activities and the Earth-to-spacecraft line of sight. In consequence, the navigation team had to rely upon the spacecraft's Doppler shift to measure the thrust in this plane. These problems were compounded by the fact that the primary component of the thrust was also perpendicular to the spacecraft's flight path. Changes had to be measured with respect to the craft's original velocity along that plane. These measurement problems stemmed from a navigation strategy that depended on the Earth-based, Deep Space Network to track the Mars Climate Orbiter. A number of alternative technologies might have been used. For instance, the Polar Lander mission also recruited a measurement technique known as 'Near Simultaneous Tracking'. These alternatives were not implemented or were not operational when the Climate Orbits reached the point of Mars Orbital Insertion [571]. It is important to note, however, that even if they had been implemented they may actually have contributed to the existing confusion about navigation data:

"The use of supplemental tracking data types to enhance or increase the accuracy of the Mars Polar Lander navigation solutions was discussed. One data type listed in the Mars Polar Lander Mission Planning Databook as a requirement to meet the Entry Descent Landing (EDL) target condition to a performance of better than 95 percent is the Near Simultaneous Tracking (NST). Additional data types discussed were the use of a three-way measurement and a difference range process. These data types would be used independently to assess the two-way coherent measurement data types (range and Doppler) baselined by the prime operations navigation team. During the presentations to the Mishap Investigation Board, it was stated that the Mars Polar Lander navigation team lead would be involved in the detailed analysis of the NST data. The application of a NST data type is relatively new to the Mars Polar Lander mission navigation procedure. These data types have not been previously used for Mars Climate Orbiter or Mars Polar Lander navigation. The results of the new data types in addition to range and Doppler only-solutions could potentially add to the uncertainty of the best estimate of the trajectory at the EDL conditions." [565]

Figure 10.12 introduces these technological issues into previous ECF diagrams. This diagram includes an event labelled Decision not to implement alternative tracking techniques and a condition Reliance on Doppler shift measurements and the Deep Space network exacerbated attempts to directly observe the impact of AMD events. As can be seen, this reliance upon a particular tracking technology contributed to the failure of the people-based barriers mentioned in previous sections. This analysis raises a number of additional meta-level points that can be made about the use of barrier analysis to drive the development of ECF charts. It introduces a new event into the primary sequence. This denotes the decision not to initiate the TCM-5 maneuver. Although we have distinguished between the people, process and technology-based barriers, incidents often stem from complex interactions between these different protection mechanisms. A failure in one area of a system, as we have often seen, will compromise other forms of protection. The difficulties of making direct observations about the AMD events frustrated attempts to quantify any residual navigation error. The significance of any such error was not fully understood; key personnel were not familiar with the Climate Orbiter's operating characteristics.

Previous paragraphs have used a relatively high-level barrier analysis to refine and guide the development of more detailed ECF charts. For example, Table 10.1 is relatively abstract when compared with the more detailed events and conditions in Figure 10.12. It is, however, possible to construct barrier tables that capture more detailed observations about the problems that exacerbate mission failures. Table 10.2 builds upon the previous analysis to look at the more detailed reasons why the software bugs in the trajectory modelling were propagated beyond the development of the

Climate Orbiter. These reasons focus on three potential barriers. The Software Interface Specification describe the units that were to be used within the project. In order to understand the failure of the Climate Orbiter, it is important to understand why this specification was not followed. The development and operations team also had detailed plans for the validation of system components. Again, it is important to understand why these plans failed to ensure the success of the mission. Finally, JPL supported a form of incident reporting system known as the Incident, Surprise, Anomaly scheme. This was deliberately intended to ensure that concerns, such as the anomalous data from the ground navigation software, was not ignored. If it had been reported to the system, there is a good chance that the concerns of the navigation team would have been addressed before TCM-5.

Hazard:	Target:
Impact/Re-Entry	Mars Climate Orbiter
Level 2 Technology: Incorrect Trajectory Modelling	
Barrier	Reason for failure?
Software Interface Specification	No software audit to ensure SIS conformance
	Poor navigation-spacecraft team communication.
	Inadequate training on importance of SIS
Software Testing and Validation	Unclear if independent tests conducted.
	Failure to recognise mission critical software.
	Poor understanding of interface issues
Incident Reporting Systems	Team member did not use ISA scheme.
	Leaders fail to encourage reporting.
	Domain experts not consulted.

Table 10.2: Level 2 Barrier Table for the Loss of the Climate Orbiter.

The Mars Surveyor Operators Project was guided by a Software Interface Specification (SIS) that both the format and units of the AMD file. This file was generated by SM_FORCES software running on ground-based computers. In order to satisfy the SIS requirements it was anticipated that this software would use metric units of Newtons per second to represent thruster performance data. As we have seen, however, the SM_FORCES software used English units of pounds per second. Subsequent processing of the AMD data by the navigation software algorithms therefore, underestimated the effect of AMD events on the spacecraft trajectory. The data was incorrect by a factor of 4.45; the ratio of force in pounds to Newtons. The SIS was intended to provide an important barrier against the type of software problems that led to the navigation software error. The previous analysis does not, however, explain why the SIS failed to protect the system in the manner intended. Primary and secondary investigations identified inadequate training a key reason why development engineers failed to satisfy the interface requirements: “the small forces software development team needed additional training in the ground software development process and in the use and importance of following the Mission Operations SIS” [565].

Inadequate training about the importance of the SIS was compounded by a lack of training about appropriate testing techniques for the ‘small forces’ software. Not only did this increase the likelihood that the software would not comply with project interface requirements but it also reduced the likelihood that any anomalies would be identified. The investigators expressed a number of additional concerns about the testing procedures that were used during the development of the Climate Orbiter. It was unclear whether or not the ground software had been inspected by an independent validator. This lack of rigour can be explained by a possible perception that the small forces software was not ‘mission critical’. It can, therefore, be argued that the technological defences of an independent verification and validation program were breached by a managerial lack of oversight and the decision not to perform a system level hazard analysis.

The Mishap Board recommended that the Polar Lander teams should develop a verification matrix. One axis would denote all mission-critical project requirements. A second axis would denote the subsequent ‘mile-posts’ in mission development. A cell in the table would only be ticked

if developers could present test results to demonstrate that the associated requirement had been met. The intention was that the verification matrix would explicitly record the test results for various requirements in Interface Control Documents, such as the SIS. It was also argued that the technical end-users of ground software applications should be required to sign-off these verification matrices.

Previous paragraphs have argued that limited training of key development staff led to an ignorance about the SIS and to inadequate testing of ground based software, including the small forces routines. Inadequate training also compromised a number of other barriers that might have protected the Climate Orbiter. In particular, the secondary investigation found members of the project team that did not understand the purpose or mechanisms of the Incident, Surprise, Anomaly (ISA) scheme. This finding is particularly important given the topic of this book. The ISA system was the primary means of providing information about adverse occurrences. Potential faults were logged with the system. Any subsequent remedial actions were then carefully monitored to ensure that the underlying issues were dealt with:

“A critical deficiency in Mars Climate Orbiter project management was the lack of discipline in reporting problems and insufficient follow-up. The primary, structured problem-reporting procedure used by the Jet Propulsion Laboratory the Incident, Surprise, Anomaly process was not embraced by the whole team. Project leadership did not instill the necessary sense of authority and responsibility in workers that would have spurred them to broadcast problems they detected so those problems might be articulated, interpreted and elevated to the highest appropriate level, until resolved.” [571]

It is difficult to underestimate the importance of these points. If the navigation anomalies has been reported to the ISA system then there is a good chance that the navigation and spacecraft operations teams would have been requested to provide a coordinated response. This response might also have involved mission scientists who had the most knowledge of Mars, of the on-board instruments and of the mission science objectives. The investigators subsequently argued that their input could well have reversed the decision not to perform the TCM-5 maneuver.

Figure 10.13 presents an ECF chart that captures some of the more detailed events and conditions that helped to undermine the defences against software ‘bugs’ on the Climate Orbiter mission. As can be seen, the insights provided by the previous barrier analysis relate to two different stages in the mission. The top-left of the diagram represents the developers’ failure to use the SIS or then to discover that this interface had been violated. Events have been introduced to represent that the SM_Forces routines are written using imperial and not metric units for thruster performance and that Limited independent testing of the ground based SM_Forces routines took place. In contrast, the lower left-hand side of Figure 10.13 represents the failure of the operational staff to report the apparent navigation anomaly using the ISA scheme.

As can be seen, training failures are represented by conditions in both areas of this diagram. This observation has a more general significance beyond our analysis of the Climate Orbiter mission. Chapter 2.3 argued that training is often perceived to be a low cost work-around for a range of deeper design, development and management problems. It should not, therefore, be surprising if inadequate training is often identified in the role of a failed barrier or inadequate form of protection. It is regrettable that ‘improved training’ is often advocated as the remedy for this problem. More might be gained from a closer examination of why training failed to provide necessary protection in the first place.

10.2.3 Change Analysis

Previous section have shown how barrier analysis can direct the construction of ECF diagrams. Previous sections have not, however, shown that ECF diagrams can be used to distinguish between root causes and contributory factors. This is a deliberate decision. As we shall see, investigators must consider a range of information about the course of an incident before attempting such a causal analysis. The following paragraphs, therefore, present a further techniques that can be used to identify further information that can be used to identify the root causes of an incident. Rather

than repeat a barrier analysis for the Polar Lander incident, this section shows how change analysis can also be used as a precursor to this causal interpretation of an adverse occurrence.

The US Department of Energy [208], Occupational Safety and Health Administration (OSHA) [651] and NASA [572] all advocate change analysis as a key analytical tool for incident investigation. This technique was pioneered by Johnson in the year immediate after the Second World War. It was then developed for use by the US Airforce by Kepner and Tregoe in the Rand Corporation [250]. Change analysis can be used to determine whether or not abnormal working practices contributed to the causes of an adverse occurrence. The focus of this analytical technique is justified by the observation that deviations from normal operations are often cited as a cause in many accidents and incidents [208]. It is important to emphasise, however, that these changes are often made with the best intentions. For instance, new working practices may help to ensure that organisations satisfy regulatory requirements. Alternatively, new production processes can be introduced to improve organisational efficiency. Problems arise not from the intention behind such changes but from the difficulty of predicting the impact that even small changes can have upon the operation of complex, technological systems. Even apparently beneficial changes can have unintended consequences that, in the medium or long term, can help to produce incidents and accidents.

In incident investigation, change analysis can be applied to identify the differences between what was expected to occur and what actually did occur during. OSHA's guidelines for incident and accident investigation include a brief tutorial on change analysis [651]. The following list enumerates the key stages in the OSHA approach. The US Department of Energy omit the final two stages and, instead, argue that investigators should feed the results of any change analysis into techniques that are intended to distinguish root causes from contributory factors [208]. They recommend that these findings should inform the development of the ECF charts, introduced in this chapter:

1. Define the problem.
2. Establish what should have happened?
3. Identify, locate and describe the change.
4. Specify what was and what was not affected.
5. Identify the distinctive features of the change.
6. List the possible causes.
7. Select the most likely causes.

Both the Department of Energy and OSHA provide relatively high-level guidelines for the application of change analysis. This is important because they provide investigators with an overview of the key stages that contribute to this technique. Unfortunately, these high-level summaries can also hide some of the underlying problems that complicate change analysis within many incident investigations. For instance, it is not always easy to determine what ought to happen during normal operation. The Polar Lander and Climate Orbiter missions had many unique characteristics that made them very different from similar projects. On the other hand, it is unclear whether or not it is possible to define what might be expected to happen during a normal NASA mission. The pressure to use leading-edge technology in pursuit of heterogeneous scientific objectives makes each mission very different from the last. Even in systems that have a greater 'routine', it can be difficult to identify operating norms. For example, the Department of Energy guidelines suggest that investigators use blueprints, equipment description documents, drawings and schematics, operating and maintenance procedures, job/hazard analyses, performance indicators etc to determine the nominal operating conditions before any incident [208]. However, subtle differences often distinguish the ways in which different plants operate the same process. Even within a plant, there will be differences in the performance of different shifts and of individuals within those shifts. Similarly, the notion of an accident-free or ideal situation can be difficult to sustain in many industries. For instance, some oil installations operate running maintenance programs. Temporary fixes are used to resolve non-critical failures. This enables operations to continue until a scheduled maintenance period.

This interval is used to conduct longer-term repairs. Such maintenance schemes raise a number of questions about what is, and what is not, a nominal state. For instance, operators view the system as operating normally even though it requires longer-term maintenance. This may seem to be an isolated example. This argument can, however, be applied to a more general class of systems. Most applications continue to operate in spite of documented failures in non-critical components. Some authors have gone further and argue that complex, safety-critical systems are unlikely to be error-free [677]. They always involve adaptations and work-arounds because it is impossible for designers and operators to predict the impact that the environment will have upon their systems.

Further problems stem from the effects of compound changes. For example, operating practices and procedures evolve slowly over time so that official documents may reflect a situation that held several years previously. Under such circumstances, previous distinctions between normal and abnormal practices can become extremely blurred. Other problems arise when changes that occurred several years before are compounded by more recent changes. The change analysis guidelines suggest that investigators should address such situations by developing several baseline or nominal situations. The events during an incident should be contrasted with normal working practices immediately prior to any failure and also with normal working practices in the years before to any previous change:

“...decreases in funding levels for safety training and equipment may incrementally erode safety. Compare the accident scenario to more than one baseline situation, for example one year ago and five years ago, then comparing the one and five year baselines with each other can help identify the compounding effects of change.” [208]

Chapters 5.4 and 6.4 have already described the difficulties that can arise when investigators must piece together the events that contribute to a particular incident. Automatic logging systems can be unreliable and seldom capture all critical aspects of an adverse occurrence. It can also be difficult to interpret the information that they do capture. Individuals may be unable to recall what happened in the aftermath of an adverse occurrence. In the aftermath of an incident, there is also a temptation for operators to describe violations as abnormal occurrences even though they may have formed part of everyday working practices. Organisation, managerial and social pressures influence their participation in a primary and secondary investigation. Inconsistencies, omissions and ambiguity are a continual problem when investigators must form coherent accounts from eye-witness statements. All of these factors combine to frustrate attempts to determine ways in which an incident differed from ‘normal’ practice. Change analysis must also consider a number of further issues. It is usually insufficient simply to contrast normal behaviour with the abnormal events that occur during an incident. Once an incident has occurred, it is also important for investigators to determine the success or failure of any remedial or mitigating actions. Given that an incident occurred, it is important to determine whether or not the response followed pre-determined procedures.

These caveat are important because they identify some of the practical difficulties that emerge during the application of change analysis. It is also important to notice, however, that they do not simply affect this analytical technique. The problems of eliciting evidence and reconstructing an incident are common to all incident investigation. Change analysis is unusual because it forces investigators to explicitly address these issues during their analysis. Other techniques, including barrier analysis, make no distinction between the normal and abnormal events that contribute to an incident.

Meta-Level Change Analysis

Reason [702] argues that incidents and accidents often stem from underlying changes in the structure of complex organisations. Change analysis can, therefore, begin in a top-down fashion by considering the organisational context in which the Polar Lander mission took place. In particular, it is important to consider the consequences of the “Faster, Better, Cheaper” strategy that was introduced by the NASA Administrator, Daniel Goldin. He assumed command at a time of shrinking financial resources caused by the recession of the early 1990’s. The US government had responded to global economic problems with a program of deficit reduction that affected many including education, healthcare and housing. Golding was faced by a situation in which NASA was likely to receive insufficient funds to

cover all of its future programme commitments. He, therefore, conducted a thorough review of both existing and future projects using ‘red’ and ‘blue’ teams. These groups were to analyse both the programmes themselves and their organisational context. Blue teams examined their own programs for creative ways to reduce cost without compromising safety or science. Red teams were composed of external assessors who were intended to bring in new ideas and to ensure that those ideas were realised. This review began in May 1992 and had an almost immediate impact. By December 1992, it was claimed to have delivered a seventeen percent reduction in costs [577].

The cost improvements and efficiencies that were achieved under the new “Faster, Better, Cheaper” initiative had a profound impact on the relationship between NASA and its contractors. As we shall see, changes in this relationship were at the heart of the problems experienced during the Climate Orbiter and the Polar Lander missions. In particular, an Independent Cost Assessment Group was set up to ensure that cost estimates were as accurate as possible. This followed a General Accounting Office report into a sample of 29 NASA programs that identified an average cost growth of 75 percent. Goldin argued that “We can not tolerate contracts so fluid, that the product we bargained for in no way resembles what we end up with... We are partners with industry, but we will hold you [contractors] accountable for what you sign up to deliver and ourselves accountable for establishing firm requirements” [578].

It is difficult to find a precise definition of what the “Faster, Better, Cheaper” initiative was supposed to imply at a project level. The Mars Program Independent Assessment Team was formed after the loss of the Polar Lander [570], it identified the following components of this initiative:

- *Create smaller spacecraft for more frequent missions.* The creation of smaller, more frequent missions was intended to increase the opportunities for scientists, and the public, to participate in NASA’s work. This approach was also perceived to have the additional benefit of distributing risk across the increased number of projects. The “Faster, Better, Cheaper” strategy distributes the risk of achieving science objectives among more missions thus minimising the impact of a single mission failure;
- *Reduce the cycle time throughout a project.* Increased mission frequency was intended to help introduce scientific and engineering innovations. This would be achieved by reducing project lead time. Such reductions were not to be made by the arbitrary curtailment of development or implementation time. They were to be achieved by the elimination of inefficient or redundant processes and, especially, through the use of improved management techniques and engineering tools. In the Polar Lander and Climate Orbiter missions, this involved greater responsibilities for line management within individual project contractors;
- *Use new technology.* The “Faster, Better, Cheaper” strategy relied upon the integration of new technology into many different aspects of each mission. New technology was intended both to increase the scientific return of each mission, to reduce spacecraft size and to limit overall mission cost. It was, however, recognised that new technologies must “be adequately mature” before being incorporated in a flight program [570]. This use of innovative technology was also intended to increase public interest in NASA programs;
- *Accept prudent risk if they are warranted by the potential rewards.* It was recognised from its inception that the “Faster, Better, Cheaper” implied taking risks; “in all cases, risks should be evaluated and weighed against the expected return and acknowledged at all levels” [570]. Rather than using flight-proven techniques, programs were encouraged to incorporate new technologies if they showed promise of significantly increasing mission capabilities or improving efficiency. The use of the term ‘prudent’ in many of the “Faster, Better, Cheaper” documents was intended to ensure that these technologies underwent a rigorous testing and validation prior to their use in flights. This was encapsulated in the maxim ‘Test-As-You-Fly/Fly-As-You-Test’; validation should provide a close approximation of the eventual mission characteristics.
- *Use proven engineering and management practices to maximise the likelihood of mission success.* The technological risks associated with this new strategy were to be addressed using proven engineering and management techniques. These techniques were to include hazard

analysis, using Fault Tree Analysis or Failure Effects and Criticality Analysis. There was an explicit concern to prevent any ‘single human mistake causing mission failure’ [570]. These established techniques were also to establish a chain of responsibilities and reporting within each project. Projects were to be reviewed by independent experts from outside the projects or implementing institutions. These individuals were to provide an overall project assessment and to review any associated risks.

This description of the “Faster, Better, Cheaper” strategy acts as a statement of what was intended by Administrator Goldin’s initiatives. It, therefore, provides an ideal or standard against which to compare the particular characteristics of the Polar Lander project. This is important given the specialised nature of such missions, change analysis has most often been applied to process industries that follow more regular patterns of production. Table 10.3, therefore, uses this approach to assess the differences between the intended objectives of the “Faster, Better, Cheaper” strategy and what went on during the Mars Surveyor’98 projects. In particular, it summarises the investigators argument that the Polar Lander team were forced to:

“Reduce the cost of implementing flight projects in response to severe and unprecedented technical and fiscal constraints... One lesson that should not be learned is to reject out of hand all the management and implementation approaches used by these projects to operate within constraints that, in hindsight, were not realistic.” [580]

It is important to emphasise that Table 10.3 does *not* compare the Polar Lander mission with missions that took place before the Goldin initiative. Such a comparison would be academically interesting but might also ignore the changing financial circumstances that have fundamentally changed the way that NASA operates in recent years.

Prior/Ideal condition	Present Condition	Effects of change
Faster, better, cheaper strategy required sufficient investment to validate high-risk technologies before launch	Mars Surveyor’98 faces pressures to push boundaries of technology and cost	Greater development effort
		Use off-the-shelf hardware and inherited designs as much as possible.
		Use analysis and modeling as cheaper alternatives to system test and validation.
		Limit changes to those required to correct known problems; resist changes that do not manifestly contribute to mission success.

Table 10.3: High-Level Change Table for the MPL Mission.

The first entry in Table 10.3, therefore, summarises the intended effects of the “Faster, Better, Cheaper” strategy on the Polar Lander mission. In contrast, NASA’s investigators found evidence to suggest that the Mars Surveyor projects pushed the limits of what was possible both technologically and within available budgets. The pressure to push the technological boundaries are illustrated by the Deep Space 2 probes. These were designed to test ten high-risk, high-payoff technologies as part of NASA’s New Millennium Program. They were to demonstrate that miniaturised components could

be delivered to the surface of another planet and could be used to conduct science experiments. The risks associated with this new technology were assessed and approved by JPL and NASA management [580]. The risk-assessment was, however, performed on the assumption that there would be a ground-based system-level, high-impact test. This test was not conducted because of budgetary constraints. Although this is a specific example, it supports the higher level observation in Table 10.3 that the Surveyor projects pushed the boundaries both of technology and cost. A further illustration can be provided by a comparison between the Mars Surveyor'98 missions and the previous Pathfinder project. Pathfinder demonstrated the successful application of a comparable range of technological innovation under the “Faster, Better, Cheaper” strategy. NASA have, however, estimated that the Mars Surveyor missions were underfunded by up to 30% in comparison with the Pathfinder [570]. This estimate is supported by the funding summary in Table 10.4.

	Pathfinder	Mars Surveyor'98 (MCO and MPL)
Project Management	11	5
Mission Engineering and Operations Development	10	6
Flight System	134	133
Science and Instrument Development	14	37
Rover	25	0
Other	2	7
Total	196	188

Table 10.4: Comparison of the Development Costs for the Pathfinder and Mars Surveyor'98 (in \$ Millions at 1999 prices).

Table 10.3 summarises the impact that budgetary pressures had upon the technological development of the Polar Lander. Developers made a number of decisions that were based on budgetary considerations but which ultimately had a critical effect upon systems engineering. These included decisions to use off-the-shelf components and inherited designs as much as possible. Analysis and modeling were also to be used as lower-cost alternatives to system test and validation. Changes were to be limited to those required to correct known problems. There was pressure to resist changes that did not directly contribute to mission success. The following sections look beyond these high level effects. Change analysis is used to analyse the detailed engineering and managerial impact of the Polar Lander's “Faster, Better, Cheaper” objectives. The results of this analysis are then used to inform the ECF charts that were presented in Figures 10.6 and 10.7.

In passing, it is worth noting that Table 10.3 illustrates some of the limitations of change analysis at this relatively high level of abstraction. It does not explain the reasons why the Surveyor'98 project adopted this extreme version of Goldin's policy. Subsequent investigations argued that this was due to ineffective communication between JPL management and NASA Headquarters. NASA Headquarters thought it was articulating program objectives, mission requirements, and constraints. JPL management interpreted these statements as non-negotiable program mandates that specified particular launch vehicles, costs, schedules and performance requirements [570].

Figure 10.14 illustrates the way in which the findings from an initial change analysis can be integrated into a high level ECF chart. This is a relatively straightforward process because the *present condition* in a Change Analysis, such as Table 10.3, can be directly introduced as a condition within an ECF chart. In Figure 10.14 this is denoted by the note that is labelled Mars Surveyor'98 faces pressures to push boundaries of cost and technology. The change analysis does not, however, identify which events this present condition will effect within an ECF chart. The node labelled Launch approved has, therefore, been introduced into Figure 10.14. Later sections will refine this high-level event to look at a number of specific events that were affected by the Faster, Better, Cheaper strategy. The change analysis illustrated in Table 10.3 also documented a number of

effects that stem from the higher-level pressures to innovate and cut costs. For example, previous paragraphs have mentioned the policy to exploit off-the-shelf hardware and inherited designs as much as possible. These effects cannot be introduced directly into ECF charts. As we shall see, they occasionally refer to particular events. In this instance, they denote more specific conditions that influence the events leading to the loss of the Polar Lander. This illustrates the important point that analysts must still interpret and filter the information that is obtained using techniques such as change and barrier analysis. There is not automatic translation between the information that is derived from these approaches and their graphical representation in an ECF chart.

People: Changes in Staffing Policy

One aspect of the “Faster, Better, Cheaper” strategy was that NASA was to profit by a greater involvement with commercial organisations. The intention was to retain a civil service and JPL core competency for in-house science, research and engineering. Aerospace operations, including the operation of the Space Shuttle and the Surveyor program, were to be performed by NASA contractors. There was also a plan to transfer program management responsibility to the field Centers from NASA Headquarters. The 1996 budgetary statement also included a commitment to performance-based contracting:

“\$100 million savings are presently projected as a result of implementing performance-based contracts for aeronautical research and facility maintenance and operations. The savings come from reducing contractor staffing levels by asking the contractor to use their ingenuity in carrying out the required work. NASA will specify what we want and when it is needed vs. specifically directing the contractor not only what and when, but also how to do the job. This will involve conversion of many current NASA cost-reimbursement/level-of-effort, specification-laden contracts.” [562]

As we shall see, this contractor ‘ingenuity’ helped to erode a number of important safety mechanisms in order to meet the relevant budgetary constraints. Contractor staff habitually worked excessive amounts of overtime. There was often only a single expert available within key mission areas.

Prior/Ideal condition	Present Condition	Effects of change
Greater JPL line-management involvement in the project.	LMA staff found it hard to fulfill mission requirements with available resources.	LMA used excessive overtime to complete work on schedule.
		Many key technical areas were staffed by a single individual.
		Lack of peer interaction.
		Breakdown in intergroup communications.
		Insufficient time to reflect on unintended consequences of day-to-day decisions.
		Less checks and balances normally found in JPL projects.

Table 10.5: Change Summary Table of MPL Staffing Issues.

Table 10.5 summarises the differences between the planned use of contract management and the experience of the Polar Lander mission. The intention was to reduce costs by relying on the contractor's existing management structure to run the day to day operation of the project. The ten or so JPL staff who were involved in the project were primarily intended to provide higher-level oversight. This was a departure from previous JPL projects and the result was minimal involvement by JPL technical experts.

It is worth reiterating that the project team was expected to deliver a lander onto the surface of Mars for approximately one-half of the cost of the Pathfinder mission. Under such constraints, it was difficult for the contractor's staff to meet their commitments within the available resources. LMA used excessive overtime in order to complete the work on schedule. Many development staff worked for sixty hours per week [580]. Some worked more than eighty hours per week for extended periods of time. Budgetary constraints created further technical problems because key areas were only staffed by a single individual. This removed important protection mechanisms because it became difficult to arrange the continual peer review and exchange of ideas that had characterised previous projects. The workload may also have jeopardised communications between technical disciplines. There was insufficient time and workforce available to provide the checks and balances that characterised previous JPL missions.

Figure 10.15 provides a further illustration of the way in which change analysis can be used to inform the construction of an ECF chart. As can be seen, the additional analysis of staffing issues has helped to identify a number of conditions that affected both the development and the subsequent validation of the lander's design. As a result, the higher-level conditions that were identified in Figure 10.14, such as use analysis/modelling as cheaper alternatives to direct testing, have been reorganised into the three strands shown in Figure 10.15. These strands distinguish between conditions that relate narrowly to staff limitations, such as the use of single individuals to cover key technical areas, from wider issues relating to the technological demands and validation of projects under the faster, better, cheaper strategy. This illustrates another important point about the process of integrating the findings of barrier and change analysis into ECF charts. The introduction of new information can force revisions to previous versions of the diagram. These revisions may result in conditions or events being removed, merged, edited or moved.

Figure 10.15 introduces a further extension to the ECF notation. A horizontal parenthesis is used to indicate that conditions from a high-level change analysis and an analysis of staffing issues influence both the development and the launch approval process. Subsequent analysis might avoid this additional syntax by omitting one of the first two events in this diagram. This has not been done because some conditions, such as the lack of peer interaction, may not only have affected the decision to launch but also the development process that led to that event. Alternatively this additional syntax could be omitted if conditions were assigned to either the development or the launch approval events. For example, the use of analysis and modelling rather than direct testing might be associated with the decision to launch rather than the completion of the development phase. Such distinctions seem to be arbitrary and have, therefore, been avoided.

Technology: Changes in Innovation and Risk Management

A number of consequences stemmed from these changes in the staffing of the Polar Lander project. In particular, the communications problems that were noted by the investigators may have compromised necessary hazard analysis. In order to assess the impact of this, it is again important to establish NASA policy for an 'ideal' approach to risk management:

“To reduce risk, we need to manage our projects systematically, especially if we expect to be successful with faster, better, cheaper projects. The Risk Management process efficiently identifies, analyses, plans, tracks, controls, communicates, and documents risk to increase the likelihood of achieving program/project goals. Every project should have a prioritized list of its risks at any point in the life cycle, along with the programmatic impacts. The list should indicate which risks have the highest probability, which have the highest consequences, and which need to be worked now. It means that all members of the project team should have access to the risk list so that everyone knows what the risks

are. It means that the project team members are responsible for the risks. The team should work to reduce or eliminate the risks that exist and develop contingency plans, so that we are prepared should a risk become a real problem... From the beginning of a project, the Project Manager and team should have an idea of what the ‘risk signature’ of the project will be. The risk signature will identify expected risks over the course of the project and when the project risks are expected to increase and decrease. During the project, risks should be tracked to determine if mitigation efforts are working. ” [574]

This policy is promoted through a range of publications and courses that are supported by NASA’s Office of Safety and Mission Assurance. Change analysis again provides a means of contrasting these ‘ideals’ with the experience of the Polar lander project. Table 10.6 provides a high level view of the differences that emerge.

Prior/Ideal condition	Present Condition	Effects of change
Adequate risk assessment at system level	No system-level Fault Tree analysis was formally conducted or documented	Bottom-up Failure Modes, Effects and Criticality Analysis hides higher-level interaction/systemic issues
		No risk analysis of propulsion, thermal and control interaction.
Adequate risk assessment at subsystem level	Fault-tree analysis treated inconsistently for different subsystems	Bug in timer for up-link loss found in Fault Tree after loss of flight.
		Premature trigger of touchdown sensor found in Fault Tree before Entry, Descent and Landing but not guarded against.
Project management maintains explicit risk-signature for the project	No risk assessment for going beyond Preliminary Design Review with 15% mass margin.	Management focus on mass reduction not risk reduction activities.

Table 10.6: Change Summary Table of MPL Risk Management.

This table suggests that risk analysis should have been conducted in a systematic manner across the various subsystems but also at a project level. There was no explicit attempt to model the way in which system-level, mission, risks changed over time. NASA refers to this model as the risk signature of a project [580]. It is important because it provides managers with a means of tracking how particular development decisions can affect the risk-margins that are eroded by particular development decisions. For instance, the preliminary design review decided to proceed with only a 15% margin between the predicted mass of the Polar Lander and the capabilities of the chosen launch vehicle. This mass assessment also failed to account for a number of outstanding mass commitments. Previous projects might have anticipated a mass margin of at least 25%. This events illustrate how key decisions were informed by cursory risk assessments. The decision to proceed with a 15% mass margin also had a significant impact upon subsequent risk management. Project resources were diverted into mass reduction rather than risk reduction activities [580].

Failure Modes, Effects and Criticality Analysis (FMECA) was used to support many areas of systems engineering. This technique is, however, driven by a bottom-up analysis of failure modes. It cannot easily be used to analyse the interactions between complex sub-systems. System level properties are often lost when FMECA is used to analyse the failure modes of complex systems. Top-down risk analysis techniques can be used to overcome these limitations. A Fault Tree analysis was, therefore, conducted for specific mechanisms and deployment systems. This analysis was only conducted for those systems that were perceived to be particularly vulnerable, for instance, because they lacked any form of redundancy. As mentioned, there was no evidence of any system level fault tree analysis. In particular, there was an ‘incomplete’ analysis of the hazards that might emerge from the interaction between propulsion, thermal and control systems [580].

The problems of risk management not only affected the risk signature of the project and the hazards associated with subsystem interaction, further problems also affected individual subsystems. For example, there was a problem in the software that was designed to automatically re-establish communications links if the up-link was lost during the Entry, Descent and Landing phase. This bug was not detected before launch or during the cruise phase of the flight. A Fault Tree analysis identified this as a possible failure mode after the Polar Lander had been lost. This led to a more detailed examination of the code. External reviewers were then used to validate the hypothesised failure. Even when risk management techniques did succeed in identifying a potential failure mode, sufficient actions were not always taken to ensure that the hazard could not arise. The Mission Safety and Success Team performed a fault-tree analysis of the Entry, Descent and Landing stage. The team then conducted an analysis to determine whether or not the design afforded sufficient protection against the identified hazard. They identified a potential failure if the Hall effect sensors received premature touchdown signals. This scenario is represented in Figure 10.7. They were, however, satisfied by the software design and testing that was provided by the contractors.

Figure 10.16 incorporates the insights from Table 10.6 into an ECF chart. The change analysis helps to identify some of the conditions that influenced events leading up to the loss of the Polar Lander. As before, some of these conditions affected many different aspects of the development process. These include the lack of any system level fault tree and the inconsistent way in which hazard analysis was performed within individual subsystems. Figure 10.16 also illustrates the way in which change analysis can be used at a more detailed level to assess the impact that departures from ‘expected practice’ had upon particular events. In particular, the lack of any assessment of the risks associated with proceeding on a mass margin of only 15% had a knock-on effect when management spent increasing amounts of time on mass reduction rather than risk mitigation. These two conditions are associated with the Preliminary Design Review. This event marks a critical stage when the projects mass margins are first established.

It is important to note that Figure 10.16 illustrates some of the limitations of the ECF notation. For example, the lack of any risk assessment for the 15% mass margins is associated with the Preliminary Design Review. This condition had knock-on effects that influence many subsequent events. In particular, the managerial focus on mass reduction is shown in Figure 10.16 as affecting the Preliminary Design Review. It also clearly affected subsequent risk assessments. Unfortunately, this is difficult to denote within the existing ECF syntax. Such limitations have inspired researchers to investigate a host of more ‘advanced’ techniques. Some of these have been introduced in Chapter 8.3. It is, however, important to note the complexity of the situation that is being analysed. A condition, the lack of any risk analysis for the 15% margin, influenced an event, the Preliminary Design Review. The consequences of this event, and in particular the decision to proceed with a 15% margin, imposed conditions upon the rest of the development process, managers had to focus on mass reduction. Such situations could be denoted within the existing ECF syntax. Edges might be drawn between conditions and events that occur later in an incident sequence. This would, however, result in a proliferation of interconnections between conditions and events. Alternatively, a cross-referencing scheme might be introduced so that conditions could be repeated at different points within an ECF chart. It is worth emphasising that most analytical techniques suffer from similar problems. The process of scaling-up from small scale studies often leads to a point at which the notation fails to capture important properties of an incident. These problems can usually be addressed through accretions to the syntax and semantics of the notation. Unfortunately, this leads to problems in

training others to use the new hybrid technique. This is a serious problem. Such notation extensions can only be justified if they provide benefits to ‘real-world’ incident investigators. Many notations have been developed and extended without any practical validation.

Previous sections have focussed on high-level changes in the way in which the Polar Lander mission was managed. In contrast, Table 10.7 assesses the impact of particular technological decisions. It is important to emphasise, however, that many of these decisions were motivated by higher-level management objectives. It is also important to emphasise that these objectives were extremely complex and, potentially, contradictory. On the one hand, budgetary constraints made it essential for NASA to justify its expenditure on technological innovation. On the other hand, many previous missions exhibited an understandable conservatism based on the feeling that mission success could be assured through the use of proven technology. This conflicts can be clearly seen in the Federal review of NASA laboratories. This formed part of President Clinton’s wider initiative that also examined the Department of Defence and Energy’s facilities. The resulting report argued that NASA’s relatively large scientific research budget produced “limited opportunities for developing technologies” to address the faster, better, cheaper strategy [573]. They also acknowledged, however, that the gap between technology development and technology utilization was the most significant problem faced by NASA’s Space Technology Enterprise. The review also reported the strong tendency within NASA to incorporate only “flight-proven technology” into space-flight missions.

These diverse factors created unusual effects on the Polar Lander project. On the one hand, the Deep Space 2 project shows a strong desire to assess the capabilities of a range of technological innovation. On the other hand, the Lander itself was developed with the explicit intention of borrowing as much as possible from previously successful mission. The Polar Lander was equipped with a disk-gap-band parachute that was identical to the one used on the Pathfinder mission, except that the Pathfinder logo had been removed. It also used an Eagle-Picher type of battery from the same batch as the one used on Pathfinder. This overall policy was, however, compromised when developers identified potential opportunities to reduce the project budget. For example, the lander exploited off-the-shelf engines that forced revisions to the initial configuration. Such technical innovations met the objectives espoused by the proponents of faster, better, cheaper. They also increased the level of uncertainty associated with the Lander’s eventual performance.

Prior/Ideal condition	Present Condition	Effects of change
Throttle valve for descent engines.	Pulse-mode control.	More difficult terminal descent guidance algorithm.
Lander design based on 2 canted engines in 3 locations.	4 smaller off the shelf engines in 3 locations.	Additional design and validation complexity.
Entry, descent and landing telemetry is available	Entry, descent and landing telemetry was not available	Problems in determining causes of mishap to inform future of program.
Downlink possible through omni-antenna	X-band down-link dependent upon MGA being pointed accurately at Earth.	Reduced chance of obtaining engineering data after anomalous landing.

Table 10.7: Change Summary Table of MPL Technological Issues.

As mentioned, Table 10.7 summarises the consequences of pressures to exploit technological innovation as a means of supporting the faster, better, cheaper strategy. This assessment is supported by the NASA investigators. The investigators found that the decision not to have EDL telemetry was defensible in terms of the project budget. It was, however, indefensible in terms of the overall program because it placed severe constraints on the amount of information that could be gleaned

from any potential failure. Finally, communications were compromised by the decision to base the Lander's X-band down-link on a medium gain antenna that had to be accurately pointed at the earth. There was no X-band down-link through the more 'forgiving' omni-antenna. This "reduced the ability to get health and safety engineering data in an anomalous landed configuration. [580]". The decision to use pulse-mode control for the descent engines avoided the cost and risk of qualifying a throttle valve. This, however, increased the complexity of the descent guidance algorithm and introduced further risks into the propulsion, mechanical, and control subsystems. The lander configuration required at least two canted engines in each of three locations for stability and control. The project elected to use four smaller off-the-shelf engines at each location.

Figure 10.17 again shows how the findings of a change analysis can be integrated into an ECF chart. In particular, this diagram focuses on the communications issues that restricted communication both during and immediately after the Entry, Descent and Landing phase of the mission. Table 10.7 captured the observation that, in retrospect, it would have been better to have provided telemetry data during Entry, Descent and Landing. The decision not to provide this facility was justified by the argument that "no resources would be expended on efforts that did not directly contribute to landing safely on the surface of Mars" [580]. As can be seen, Figure 10.17 represents this analysis as two conditions labelled *Entry, descent and landing telemetry is not available* and *Problems in determining cause of mishap* make it hard to identify lessons for future systems. These conditions are, in turn, linked to previous ECF charts by introducing an event that represents the establishment of the mishap board. Their work was complicated by the lack of telemetry data.

Figure 10.17 also includes conditions that represent the potential effects of a communication failure. This is done by the conditions that are labelled *X-band down-link is dependent upon medium gain antenna being accurately pointed at Earth* and *Reduced chance of obtaining engineering data after anomalous landing*. This raises a further problem in the application of ECF charts as a means of modelling complex incidents and accidents. Previous sections have mentioned that the lack of any telemetry data makes it difficult for investigators to be certain about the exact causes of the failure. In consequence, Figure 10.17 represents a scenario in which the Lander is lost through the software bug in the handling of spurious signals from the Hall effect sensors and the Deep Space 2 probes are lost from electrical failures at impact. If, however, the software bug did lead to the loss of the lander then the decision to rely on the Medium Gain Antenna for the X-band up-link becomes of secondary importance to this incident. The chances of the Lander surviving the resultant impact with the planet surface are so remote that it this decision would have had little effect on the incident. Figure 10.17, therefore, introduces a double-headed line to illustrate that the X-band link may be significant for other failure scenarios or for future missions but that it is of limited relevance to this incident.

Table 10.7 also summarises the inspectors argument that the limited budget created a number of problems in assessing the cost-risk tradeoff for particular technological decisions. The difficulty of making such an assessment led to unanticipated design complexity. The decision to use pulse-mode control for the descent engines avoided the cost and risk of qualifying a throttle valve. This, however, increased the complexity of the descent guidance algorithm and introduced further risks into the propulsion, mechanical, and control subsystems. The lander configuration required at least two canted engines in each of three locations for stability and control. The project elected to use four smaller off-the-shelf engines at each location. Figure 10.18 represent two events in the development of the Lander: *Decision to use pulse mode control* and *Decision to use off-the-shelf engines in 4x3 configuration*. These events provide a specific example of the way in which technological innovation and cost constraints often demand increased development effort.

It is important to reflect on the process that we have been following over the last few pages. The US Department of Energy recommends change analysis as a means of supplementing an initial ECF chart. The intention is to ensure that investigation consider a range of key events and the conditions that influence those events before any causal analysis is attempted. This approach is also recommended by the NASA guidelines for 'Mishap Reporting, Investigating and Record-keeping' [572] The Polar Lander case study illustrates a number of benefits that can be obtained from this complementary approach. In particular, the change analysis provides a good means of identifying the wider contextual issues that can often be overlooked by more event-based approaches. This

is illustrated by the way in which change analysis helps to focus on the impact of managerial and organisational strategy. Our analysis has also indicated a number of potential weaknesses in the use of change analysis to inform the construction of ECF charts. Figure 10.18 only presents a small portion of the overall diagram. In ‘bespoke’ projects such as the Polar Orbiter mission, change analysis is likely to identify a vast range of potential differences from previous projects. It is important to reiterate that our case studies were deliberately chosen with this in mind, previous examples of ECF charts focus on the more routine analysis of incidents within the process industries [210].

Process: Changes in Development Practices and Reviews

Previous sections have identified differences between recommended risk management practices and the approach that characterised the Polar Lander’s development. Many of the deficiencies can be explained by resource constraints. Others can be justified in terms of the practical challenges that such ‘leading-edge’ projects pose for current analysis techniques. The limited nature of the risk assessment process during the Polar Lander project did, however, have a number of knock-on effects. For example, previous NASA projects were typified by an extensive use of redundancy as a means of combating potential failures. The Shuttle’s design was based on the maxim ‘fail operational/fail operational/fail-safe’. One failure and the flight can continue but two failures and the flight must be aborted [566]. Even in these applications, however, it is not practical to develop fully redundant systems. In consequence, risk analysis guides the application of redundancy to the most mission-critical areas of a design. However, the lack of any system-wide hazard analysis arguably prevented the effective use of redundancy to protect against failure during key phases of the mission. It was noted that “certain MPL mission phases and sequences provide coverage only for parameter dispersions that conservatively represent stochastic dispersions, but unnecessarily fail to acceptably handle anomalously large parameter dispersions created by unmodeled errors or other non-stochastic sources” [571]. In particular, there was no functional backup if the Entry, Descent and Landing failed to follow an ‘ideal’ sequence of events. Table 10.8 summarises these knock-on effects that a limited risk analysis had upon the development of the Polar Lander mission.

Table 10.8 represents more general concerns about the models that guided the Lander’s development. For instance, models were used to characterise the potential designs of the spacecraft as well as the environment in which it was intended to operate. Any inconsistencies, inaccuracies or omissions could have had profound consequences for the eventual success of the mission. Unfortunately, it is difficult to underestimate the complexity of constructing and validating such abstractions. Models that characterise one subsystem often influence, and are influenced by, many other subsystems. This creates considerable complexity because different aspects of a system are developed at different speeds. For example, thruster and software design lagged behind other Lander subsystems. Further problems complicated the use of predictive models. In particular, the small forces generated by the spacecraft could not be modeled to the level of accuracy that was required by the navigation plan. This called for precision navigation requirements that were incompatible with the spacecraft’s design.

Validation and verification techniques can be used to test a potential design under simulated operating conditions. The results of such tests also provide insights into the utility of any models that guide systems development. Unfortunately, results can be compromised if validation tests are based on the same incorrect assumptions that guide mission development. Systems will perform well under simulated operating conditions that have little relationship with an eventual working environment. The problems of conducting such validation exercises are compounded by the managerial issues that complicate any multi-disciplinary development. Insufficient instrumentation, an error in the thermal model and poor communication between the propulsion and thermal groups produced inaccurate results from the Lander’s thermal-vacuum tests. As a result, several design problems were not detected until after the launch. The Lander’s validation “was potentially compromised in some areas when the tests employed to develop or validate the constituent models were not of an adequate fidelity level to ensure system robustness” [580].

NASA standards recommend independent verification and validation as a means of avoiding

Prior/Ideal condition	Present Condition	Effects of change
Design is resilient beyond conservative stochastic parameter dispersions.	Design vulnerable to unmodeled errors or non-stochastic sources.	EDL Sequence fails under anomalous conditions
		No functional backup for several systems.
Spacecraft design should match mission requirements	Aspects of the design could not be modelled accurately enough for control	Small forces not accurately modelled for precision navigation.
Properly validated models should be used when testing is impossible	Some models not properly validated	Doubts over results for radar-terrain interaction.
		Doubts over dynamical control effects of pulse-mode propulsion.
Sufficient resources to assess interaction between propulsion, thermal and control subsystems	<p>Thermal and software design lags behind other subsystems requiring these inputs.</p> <p>There was an error in the thermal model used to support thermal-vacuum tests.</p> <p>Insufficient instrumentation of the thermal-vacuum tests.</p> <p>Poor communication between propulsion and thermal groups.</p>	<p>Partial evaluation of propulsion, thermal and control interaction.</p> <p>Inadequate thermal-vacuum tests.</p> <p>Problem with catalyst bed heaters had to be handled prior to entry.</p> <p>Remaining concerns over uneven propellant drain from tanks during descent.</p>
Sufficient resources to validate and verify software in landed configuration.	Flight software not subjected to 'system-level' tests.	Post-landing fault-response bugs only uncovered after mission loss.
		Touchdown sensing software untested with lander in flight configuration.

Table 10.8: Change Summary Table of MPL Process Issues.

such problems [560]. Tests are conducted by organisations that are not involved in the development process. In consequence, they are less likely to follow the assumptions that are embodied within system models. External auditors may also be slightly more resilient to the internal pressures that complicate the conduct of integration tests within complex development teams. Unfortunately, this form of testing is expensive. On a resource-limited project, it must be focussed on those areas of a mission that are considered to be of prime importance. Technical difficulties further complicate the validation of complex systems. These problems prevented developers from testing system performance during the Entry, Descent and Landing phase under the Martian gravity of 3/8g. Partly as a result of this, the touchdown sensing software was not tested with the lander in the flight configuration and the software error was not discovered during the verification and validation program.

Figure 10.19 gathers together the products of the different forms of change analysis that have been conducted up to this point. These conditions describe the impact of changes in staffing policy and risk assessment practices. They also outline the effects of wider changes in NASA project management strategy and in development practices. These conditions collectively describe the context in which the Polar Lander was developed and launched. As more information becomes available about particular events, investigators can draw upon this contextual information to identify particular conditions that influenced those events. This approach provides a number of benefits. The conditions identified by change analysis need not be immediately associated with particular events. For example, conditions can emerge from the documents and statements that are gathered during a primary investigation. It can be difficult to identify particular events that are associated with the information that is provided by these documents. For instance, statistical comparisons of different levels of funding on various projects provide important information about the wider context in which an incident occurs. It would, of course, be possible to invent an event so that these conditions could be linked into an ECF chart. In contrast, Figure 10.19 shows how these contextual conditions can be gathered together for integration into an ECF chart, if and when investigators need to provide additional information about the conditions that affect particular events. Investigators are free to determine whether or not they should be explicitly associated with more detailed events. The complexity of ECF charts such as Figure 10.17 is an important consideration here. If all of the conditions represented in Figure 10.19 were explicitly linked to the different events that they influenced then the resulting ECF chart would rapidly become intractable. The task of determining the appropriate level of detail in such diagrams, therefore, forms an important component of the wider causal analysis.

Figure 10.20 illustrates how conditions can be introduced to provide further information about the events that are already represented within an initial ECF chart. In this case, the change analysis identifies that the *touchdown sensing software is untested with the lander in flight configuration*. It also identifies the more general point that the *flight software was not subjected to a systems level test*. These conditions both provide insights on the software problem that was identified in the Hall Effect sensors. This, in turn, led to the hypothesised failure scenario in which there was a premature shut-down of the lander's engines.

This analysis identifies a number of important caveats about our use of change analysis to drive the construction of ECF charts. In developing an initial ECF chart, we already identified the scenario in which the lander's engines were cut at forty meters above the planet surface. This helps to direct the subsequent analysis towards any changes that might have contributed to such a software failure. On the one hand, this can be seen as beneficial because it guides the allocation of finite investigatory resources. On the other hand, the generation of an initial hypotheses may bias any subsequent change analysis. This is especially important where there are considerable differences between each mission or run of a production process. Rather than considering the wider range of potential changes, analysts are biased towards those that support pre-existing hypotheses. This argument supports Mackie's ideas about causal fields that were introduced in Chapter 6.4 [508]. He goes on to develop the notion of a causal field that describes the normal state of affairs prior to any incident. Investigators try to identify the causes of an incident by looking for disturbances or anomalies within the causal field. This causal field is, therefore, a subjective frame of reference that individuals use when trying to explain what has happened in a particular situation. If a cause does not manifest itself within the causal field then its influence is unlikely to be detected. These ideas

have a particular resonance in our use of change analysis. Both Table 10.19 and Figure 10.20 reflect subjective assumptions about what was ‘normal’ development practice. It was argued that sufficient resources should have been allocated to validate and verify software in landed configuration. Given that budgetary constraints affected almost every aspect of the Lander’s development, the selection of this particular conditions provides insights not only about the incident itself but also about the investigator’s causal field.

There is also a danger that the counterfactual arguments, which we have adopted, may also serve to compound the salience bias that we have described in the previous paragraph. Counterfactual reasoning encourages analysts to identify causes, which had they not occurred then the incident would not have occurred. There is a danger that this can lead to a search for ‘silver bullets’; the minimal set of events that might have avoided the incident. This ‘silver bullet’ approach ignores Mackie’s argument, introduced in Chapter 6.4 that there will be alternate ‘causal complexes’ that might lead to a future incident [508]. Mackie views a cause (in the singular) to be a non-redundant factor which forms part of a more elaborate causal complex. It is the conjunction of singular causes within the causal complex that leads to an outcome. The causal complex is sufficient for the result to occur but it is not necessary. There can be other causal complexes. By extension, the ‘silver bullet’ approach is likely to rectify singular causes within a causal complex. It is, however, likely to overlook other causal complexes that can lead to similar failures in the future. This is an abuse of counterfactual reasoning rather than a weakness of the approach itself. It is also important to distinguish between general and particular causation. A general cause is one which can be used to characterise a number of different instances of the same factor. For example, poor situation awareness is a general cause of aviation accidents. In contrast, a particular cause is an instance of a general cause and describes a specific example of this more general problem. Hence we can have both general and particular, singular causes.

In the context of our analysis, there is a danger that change and barrier analysis might be used to support the preliminary hypotheses that are identified in ECF charts without examining the wider causal complexes identified by Mackie. Any subsequent root cause analysis will, therefore, be focussed on an extremely limited model of an incident. It is essential to stress noted that these dangers do not stem from the notations themselves. They are strongly related to the way in which those notations are used within particular incident investigations. In particular, the primary means of ensuring an adequate analysis of the causal complexes behind an incident is to expect the same level of review by peer investigators as one would expect during the design of any safety-critical system. Figure 10.21 illustrates how change analysis can be used to search for causal complexes beyond those that are identified in an initial ECF chart. This introduces conditions to denote that software to switch from a failed up-link string to a backup up-link string contained a bug and that post-landing fault response bug was only uncovered after the loss of the mission. As can be seen from the double headed edge in Figure 10.21 these conditions relate to problems in the communication system that could have contributed to the loss of the mission but not if the engines had indeed been cut at forty meters from the planet surface.

The previous paragraphs have argued that some of the software flaws were not detected because it was untested with the lander in flight configuration. There are both technical and financial barriers to such tests. NASA, therefore, advocates the use of formal reviews to supplement direct testing. These meetings are intended to increase consensus and confidence about a proposed design. For instance, the NASA Standard 5001 for the ‘Structural design and test factors of safety for space-flight hardware’ states that:

“Standard criteria cannot be specified for general use in designing structures for which no verification tests are planned. Projects which propose to use the no-test approach generally must use larger factors of safety and develop project-specific criteria and rationale for review and approval by the responsible NASA Center. For spacecraft and other payloads launched on the Space Shuttle, these criteria must also be approved by the Space Shuttle Payload Safety Review Panel prior to their implementation.” [563]

Partly in response to the loss of the Climate Orbiter and the Polar Lander, NASA have recently published procedures for the ‘Management of Government Safety and Mission Assurance Surveil-

lance Functions for NASA Contracts' [569]. This identifies a continuum of oversight ranging from low intensity, periodic reviews to high intensity oversight, in which NASA managers have day-to-day involvement in the suppliers' decisionmaking processes. These different forms of oversight are coordinated through a surveillance plan that must be submitted within 30 days of any contract being accepted. The plan describes the safety and mission assurance functions that are necessary to assure that the contractor will meet project requirements. Independent agencies may be identified in this plan if they are to validate the results of any assurance functions. Surveillance plans must be revised to keep pace with changes in the contractors' operations. The plan and its revisions must be reviewed at least annually to determine whether or not it must be further revised. As mentioned, these requirements were not in place during the development of the Polar Lander. There are considerable dangers in applying standards that hold after an incident to identify deficiencies that led to any mishap. There, Table 10.9 restricts its analysis to those review activities that were recommended in documents such as [563] and [561].

Prior/Ideal condition	Present Condition	Effects of change
Subsystem Preliminary and Critical Design Reviews provide independent evaluation of key decisions	Contractors lacked necessary input from external sources	Flight System Manager chaired all subsystem reviews
		LMA staff approve closures on actions without independent technical support.
		Some actions did not adequately address concerns raised by reviews.

Table 10.9: Change Summary Table of MPL Review Issues.

The investigators found that the Polar Lander project did not have a documented review plan. It did, however, hold both formal and informal reviews. Each subsystem coordinated their own preliminary and critical design reviews. This informal approach was intended to reduce the level of bureaucracy that had been associated with assurance functions in other projects. This informal process was used to communicate concerns and generate requests for actions. Unfortunately, these subsystem reviews demonstrated varying levels of technical analysis. Some issues, such as the design of the G and H release nut, were examined in a meticulous and thorough manner. Others were not. For instance, the thermal control design interfaces were not mature enough to evaluate at propulsion systems critical design review. Had a subsequent review been scheduled then the developers might have discovered some the problems that were later experienced in flight.

A mission assurance manager tracked each review action to ensure that it was addressed by a written closure and that the closure was then approved by a relevant authority. This procedure was used to ensure that all actions and recommendations were closed prior to launch. These closures were, however, typically approved by LMA staff without any independent technical support. This need not have been a concern if some form of meta-level independent review had been conducted of these closures. As we have seen, however, budgetary constraints meant that there was minimal JPL technical support. LMA did not have their closures reviewed by Board members or by non-project LMA personnel. It was later argued that:

“This limitation on technical penetration of the action items and their closure is not typical of JPL projects and was probably an unintended consequence of project funding limitations. Rather than following the typical process of choosing board chairpersons

with technical expertise in functional areas from outside the project, the Flight System Manager was the chairperson of all the subsystem reviews.” [580].

In passing, it is worth noting that the problems of developing effective assurance procedures for contracted work has been a recurring theme in recent NASA mishap reports [576]. This, in part, explains the subsequent development of a comprehensive set of standards and policies in this area.

Figure 10.22 provides a final illustration of the use of change analysis as a means of expanding an ECF chart. In this case, several further conditions are introduced to annotate the development and review events that have been identified by previous stages of the analysis. This figure again illustrates the problems of associating conditions with individual events. Parenthesis are again used below the event line to indicate the potential scope of these conditions. As with previous diagrams, it would be possible to refine the events shown in Figure 10.22 so that conditions can be more firmly rooted to particular moments during an incident. This is a subjective decision, I chose not to do it in this analysis because it would have forced me to invent a number of arbitrary events. The available evidence was not in a format where I could have such distinctions. In general, this reflects the difficulty of representing persistent constraints within event-based notations. Time-lines suffer from similar problems and the solutions were almost identical in Chapter 8.3. This remains an area of current research. For now, it is important to realise that our integration of change analysis and ECF charts has exposed a number of limitations in the application of this analysis technique for a complex, technological failure.

Previous sections focussed on the ways in which particular aspects of the Polar Lander’s development may have contributed to the failure of this mission. In particular, we have identified instances in which this project adopted practices and procedures that differed from those advocated by senior management through published guidelines and policies. Limited funding and changes to NASA’s subcontracting practices helped to place heavy burdens upon the available staff. These burdens, together with particular skill shortages, had an adverse effect on the risk assessments that are intended to guide subsequent development. As a result, a number of technical decisions were made that could not easily be justified in retrospect. For example, the lack of telemetry during the Entry, Descent and Landing phase created considerable problems for investigators who must feed any relevant lessons into current and future projects. Further problems arose from the technical and financial barriers that prevented development teams from testing all aspects of the Polar Lander’s design. Such tests might have helped to identify potential problems that were not identified during a hazard analysis. Instead, a number of problems were discovered after the craft was in flight. Such problems also illustrate the way in which the Polar Lander’s project reviews had failed in their meta-level role of assuring mission success.

It is important to stress that the previous tables have been guided by an implicit form of change analysis that is apparent in the documents and records that were produced by the NASA investigators. In order to identify potential shortcomings that might have affected the mishap, they first had to analyse the recommended practices for similar development projects:

“NASA currently has a significant infrastructure of processes and requirements in place to enable robust program and project management, beginning with the capstone document: NASA Procedures and Guidelines 7120.5. To illustrate the sheer volume of these processes and requirements, a partial listing is provided in Appendix D. Many of these clearly have a direct bearing on mission success. This Boards review of recent project failures and successes raises questions concerning the implementation and adequacy of existing processes and requirements. If NASA’s programs and projects had implemented these processes in a disciplined manner, we might not have had the number of mission failures that have occurred in the recent past.” [580]

For example, the software component of the Lander development was covered by NASA standard NASA-STD-2100-91 (Software Documentation, [559]), by NASA-STD-2201-93 (Software Assurance, [560]), by NASA-STD-2202-93 (Software Formal Inspections, [561]) and by a draft form of NASA-STD-8719.13A (Software Safety, [564]). This illustrates an important limitation of change analysis. In an organisation as complex as NASA, it is likely that there will be a significant body of information

about recommended practices. It can be difficult or impossible for any individual to continually assess whether their project conforms to all of the available guidelines. As a result, it is likely that most projects will differ from the ideal. It can also be difficult for developers to learn more about successful practices from other projects. One means of addressing this problem is to provide developers with means of searching for appropriate guidelines and lessons learned. NASA provide a web-based interface to their standards library for this purpose. By extension, it can also be argued that same facilities ought to be available to help inspectors search for incidents in which these standards were not followed. Such tools can be used to identify emerging patterns of related failures within a database of incidents. Chapter 14.5 will describe some of these systems in more detail. In contrast, the following section goes on to show how ECF charts can be used to direct a causal analysis of the Polar Lander and Climate Orbiter case studies.

10.3 Stage 2: Causal Analysis

This section goes on to describe how a number of analytic techniques can be used to distinguish causal events from the mass of contextual events and conditions that are identified in preliminary ECF charts. In particular, ECF Analysis, Tier Diagramming and Non-compliance Analysis are used to filter the mass of information that is gathered during primary and secondary investigations.

10.3.1 Causal Factors Analysis

The Department of Energy guidelines argue that ECF charting must be conducted to a sufficient level of detail and that this depends upon *both* change *and* barrier analysis [208]. The NASA guidelines, NPG 8621.1, are ambiguous in this respect [572]. Barrier analysis appears as an item in the Mishap Board Checklist (Appendix J-3) but not in the list of recommended investigation techniques where guidance is provided on the other two complementary approaches. Irrespective of whether both analytical techniques are used to derive an ECF chart, the next stage is to analyse the resulting diagram to identify the causes of an incident. This, typically, begins with the event that immediately precedes the incident. The Department of Energy guidelines suggest that investigators must ask *would the incident have occurred without this event?* If the answer is yes then the analyst progresses to the next event; the event is assumed not to have had a significant impact on the course of the incident. However, if the answer is no then a number of further questions must be asked about the both the event and the conditions that are associated with it. This illustrates how causal factor analysis relies upon counterfactual argument.

A number of problems complicate this first stage of the analytical method. The first issue centres on the relationship between events and conditions. Previous sections have argued that conditions “(a) describe states or circumstances rather than happenings or occurrences and (b) are passive rather than active” [210]. Problems arise when a condition is associated with an event that is not considered to be central to the causes of an incident, i.e., the answer to the previous counterfactual question is yes. For instance, it might be argued that the Climate Orbiter might still have been lost even if more staff had transitioned from development to operations. In this case, investigators might then neglect the effect of the associated condition that the Mars Climate Orbiter is the first project for the multi-mission Mars Surveyor Operations project. It can be argued that such conditions are irrelevant because they do not directly affect the counterfactual argument that drives causal factor analysis. It can also be argued that this form of analysis places unnecessary importance on specific events and that it neglects the context in which an incident occurs. Such caveats are important because many event-based modelling techniques force investigators to invent ‘arbitrary’ events so that they can represent important elements of this context. For example, failures of omission have to be represented as negative events within an ECF line. This provides investigators with the only means of representing the conditions that influenced the omission. For example, the decision not to perform TCM-5 was influenced by the failure to understand the significance of the AMD data. This, in turn, was influenced by conditions that ranged from management changes through to a reliance on Doppler shift and the Deep Space network for tracking data. This example clearly illustrates that it is the conditions that are more important for future safety than the ‘non-event’.

Causal factor analysis is further complicated by the difficulties of applying counterfactual reasoning to complex, technological failures. For instance, how can we be sure that the Climate Orbiter would have succeeded if the Small Forces bug had been counteracted by TCM-5? There might have been other unidentified problems in the navigation software. Alternatively, TCM-5 might itself have introduced further problems. The key point here is that the previous counterfactual question refers to a particular incident. It does not ask ‘would any incident would have occurred without this event?’. Investigators cannot, typically, provide such general guarantees.

Further complications arise from multiple independent failures. These occur when an investigation reveals two or more problems that might have led to an incident. Multiple independent failures are denoted on ECF charts by different chains of events and conditions that lead to the same incident symbol. Our analysis of the Polar Lander identified two of these chains. One leads from the failure of the touchdown sensing logic. The other represents problems in the communications systems. These independent failures create problems for counterfactual arguments because the incident might still have occurred if either one of them was avoided. An investigator would answer ‘yes’ to the question ‘would the incident have occurred without the Hall Effect sensor problem?’. Conversely, they could also answer ‘yes’ to the question ‘would the incident have occurred without the communications problems after landing’. According to the ECF method they would then disregard these events and continue the analysis elsewhere! This problem can be avoided if investigators construct and maintain multiple ECF charts to represent each of these different paths. This approach has some drawbacks. For instance, it can be argued that similar events led to the touch-down sensing bugs and the software problems in the communications up-link. These common causes would then be artificially separated onto different ECF charts in order to preserve the method, described above. An alternative means of avoiding this problem is to require that investigators repeat the counterfactual question for each path that leads to an incident symbol. The question then becomes ‘would the incident have occurred in the manner described by this ECF path without this event?’.

The complex issues surrounding counterfactual reasoning about alternative hypotheses does not simply affect the Polar Lander and Climate Orbiter case studies. It is a research area in its own right. Byrne has conducted a number of preliminary studies that investigate the particular effects that characterise individual reasoning with counterfactuals [123, 124]. This work argues that deductions from counterfactual conditionals differ systematically from factual conditionals and that, by extension, deductions from counterfactual disjunctions differs systematically from factual disjunctions. This is best explained by an example. The statement that ‘the Climate Orbiter either re-entered heliocentric space or impacted with the surface’ is a factual disjunction. Byrne argues that such sentences impose additional burdens on the reader if they are to understand exactly what happened to the Climate Orbiter. In the general case, they must also determine whether both of the possible outcomes could have occurred. The statement that ‘the Climate orbiter would have re-entered heliocentric space or would have impacted with the surface’ is a counterfactual disjunction. Byrne argues that this use of the subjunctive mood not only communicates information about the possible outcome of the mission but also a presupposition that neither of these events actually took place. There has, to date, been no research to determine whether these insights from cognitive psychology can be used to explain some of the difficulties that investigators often express when attempting to construct complex counterfactual arguments about alternative scenarios. In particular, the use of counterfactual disjunctions in our analysis of the Polar Lander is specifically not intended to imply that neither actually took place. It, therefore, provides a counter-example to Byrne’s study of the everyday use of this form of argument.

Figure 10.23 presents an excerpt from the ECF chart that represents the failure of the Polar Lander mission. As can be seen, this diagram focuses on the events and conditions that may have contributed to the loss of the Deep Space 2 probes. The following paragraphs use Figure 10.23 to illustrate the application of the analytical techniques described above. In contrast to the Climate Orbiter and the Lander itself, we have not applied change or barrier analysis to this portion of the initial ECF chart. The decision to focus on this aspect of the incident is entirely intentional. The subsequent paragraphs show how causal factor analysis can be used to check whether change and barrier analysis has identified the precursors and conditions that affect the potential causes of failure. As mentioned, causal factor analysis begins with the event that immediately precedes the

incident symbol. Previous paragraphs have argued that the answer to this question is bounded by the particular ECF path that is being considered. It would, therefore, be necessary to repeat the analysis for each alternate paths leading to the same incident. Fortunately, Figure 10.23 shows a single event chain leading to the accident.

The investigator must ask whether the failure would have occurred if it was not the case that both of the DS2 probes suffer electrical failure at impact? If the answer were yes, the incident could have occurred without this failure, then the event can be classified as a contextual detail. The analysis would then move on to preceding events. In this case, however, if the electrical failure had not occurred then the probes would not have been lost. If we had omitted this event from our model, we would not have had a coherent explanation of the failure. This counterfactual argument suggests that this event is a contributory factor and that further causal factor analysis should be conducted. This causal factor analysis is based around a number of questions that are intended to ensure that analysts have identified sufficient information about key events. This information is necessary to drive any subsequent root cause analysis. It is important to stress, however, that many of the details that emerge from a causal factor analysis may already have been identified during previous stages of barrier and change analysis. This penultimate stage, therefore, provides additional assurance in the results of these other analytical techniques. The US Department of Energy guidelines argue that investigators must review the results of this analysis so that ‘nothing is overlooked and that consensus has been achieved’ [208].

Table 10.10 records the results of an initial causal factor analysis for the electrical failure event that precedes the loss of the probes shown in Figure 10.23. As can be seen, the intention behind the questions that drive the causal factor analysis is to expand on the summaries that label the ECF chart. The ECF chart is used to show *when* an event occurred. The causal factor analysis expands this to capture *what* went wrong, *why* barriers failed and *who* was involved in the event. It should be noted that these questions are a subset of those proposed by the US Department of Energy [208]. This is intended to simplify the causal factor analysis and broaden its application to include the complex, technological failures that are addressed in this chapter. It should also be noted, however, that these questions can be amended to reflect the insights that are gained during subsequent investigations. For instance, we initially had replaced *who was involved in the event?* with the question *who was responsible for the barrier?*. This original version was removed after some investigators used the answer to directly assign blame for the incident even though barriers may have been breached by a pathological conjunction of environmental behaviours and system failures.

As can be seen, the causal factor analysis in Table 10.10 helps to collate information about the development of the probes. It describes how the flight cell battery lot was delivered too late to be impact tested. Table 10.10 also includes information about validation activities. There was insufficient time to conduct a powered, fully integrated impact test on the probe communications system. Finally, it identifies groups who were responsible in approving the “proceed to launch” decision in spite of these potential concerns. These observations were not explicitly identified during previous stages in the generation of the ECF chart. They, therefore, can be interpreted as omissions that are exposed by the explicit questions in the form shown in Table 10.10. Additional events can be introduced into Figure 10.23 to represent these insights prior to the eventual root cause analysis.

The final question in Table 10.10 looks beyond the specific event that forms the focus of this analysis. In particular, it prompts the investigator to identify whether or not a particular failure forms part of a wider pattern. It follows that such annotations are likely to be revised as the causal factor analysis is repeated for many different events in an ECF chart; patterns may only emerge during the subsequent analysis. This question also provides an opportunity to explicitly identify any similarities with previous events during other incidents. Subsequent chapters will describe tools and techniques that can be used to identify common features amongst a number of different incidents. For now, however, it is sufficient to observe that primary and secondary investigations often uncover superficial similarities between the events that contribute to different incidents. These potential similarities must be investigated to determine whether or not different incidents do indeed begin to form a pattern of failure.

The causal factor analysis in Table 10.10 is untypical because we have not presented any previous barrier or change analysis to identify further events and conditions leading to the loss of the Deep

Event : Both DS2 Probes Suffer Electrical Failure at Impact	
What led to the event?	There was not enough time to conduct an impact test with a complete probe in flight configuration. Cost constraints and technical barriers also prevented such a validation.
What went wrong?	<ol style="list-style-type: none"> 1. There was no system-level impact test of a flight-like RF subsystem. Mechanical and structural validation took place at the level of brassboard and breadboard components. Many components were not electronically functional. This limited pre-test and post-test DC continuity checks. 2. The flight battery cell lot was delivered too late to be impact tested. Validation arguments were based on a preceding lot of 8 identical cells. However, one of these was physically damaged during a test but did not fail catastrophically.
How did the barriers fail?	The program exploited non-destructive tests and analytical modelling whenever possible. This was in-line with the objectives of the Faster, Better, Cheaper strategy. However, analytical models of high g impacts are unreliable and so flight qualification should have been demonstrated by tests on representative samples of flight hardware.
Who was involved in the event?	Two peer review meetings and three project level reviews established “proceed to launch” concurrence from JPL and NASA upper management. If the project team had forced an impact test for the RF subsystem and the fully integrated, powered probe then they might have missed the launch.
Is the event linked to a more general deficiency?	Many events and conditions in the Polar Lander’s ECF charts that relate to validation and review problems. The Faster, Better, Cheaper strategy is relevant to different events and conditions also.

Table 10.10: ECF Analysis of the Deep Space 2 Failure.

Space 2 mission. This was intentional because some investigations may not have the necessary resources to conduct these intermediate forms of analysis. As we have seen, it is possible to move straight from a high-level preliminary ECF chart such as Figure 10.4 to the analysis in Table 10.10. For higher consequence failures, such as the Mars Global Surveyor missions, it is likely that any causal factor analysis will build upon barrier and change analysis. Figure 10.24, therefore, integrates the events and conditions that were identified in the previous analysis of the Polar Lander incident. The relative complexity of this figure, even with the use of continuation symbols, indicates the complexity of the incident. It also provides an overview of the investigations that precede causal factor analysis.

The incident symbol in Figure 10.24 is preceded by an event, labelled *Premature shut-down of engines (40 meters above the surface)*, and by a condition, labelled *Reduced chance of obtaining engineering data after anomalous landing*. Previous sections have, however, explained that these events are mutually exclusive. This is denoted by the double-headed link between the condition and the incident symbol. If the engines had been shut-down at 40 meters then the Lander would have been destroyed on impact with the planet surface. In consequence, any problems with the communications systems are unlikely to have had a significant impact on the loss of the mission. There is a very small probability that it could have survived such an event but the NASA investigation team did not consider that it was worth pursuing. In consequence, the causal factor analysis focuses on the event that is associated with the engine shut-down.

Causal factor analysis begins by asking whether the failure would have occurred if there had not been *premature shut-down of engines (40 meters above the surface)*. The answer to this question is assumed to be no. This is the only event in the ECF chart of Figure 10.24 that leads to the loss of the mission. The enquiry process, therefore, follows the same pattern as that established for the loss of the Deep Space 2 probes. Table 10.11 summarises the answers to the questions that drive the causal factor analysis.

Table 10.10 was derived without any intermediate barrier or change analysis. In contrast, Table 10.11 benefits from the more sustained analysis described in previous sections. In consequence, the ECF prompts may simply reiterate information that was identified by the earlier forms of analysis. The premature shut-down stemmed from a spurious touchdown signal from the Hall Effect sensors. The software did not reset a variable that was set in response to this spurious signal and this ultimately indicated that the Lander had contact with the surface when it was still some 40 meters from touch-down. It is, however, likely that the causal factor analysis will prompt some novel observations. For example, Table 10.11 briefly explains how the developers were keen to balance the loading on processors during the Entry, Descent and Landing phase. This contributed to the software failure because processors sampled the Hall Effect sensors well before reaching 40 meters. The intention was to avoid any sudden processing peaks that might have been incurred by starting to poll these devices at the point at which their input was needed.

The causal factor analysis also poses some questions that were not directly addressed during previous stages in the investigation. The change analysis of the Polar Lander failure did not explicitly address the reasons *why* particular barriers failed to detect the potential bug in the landing software. As can be seen from Table 10.11, the XB0114 requirements document did not explicitly consider the possible failure modes for the landing logic. The software engineers were not informed of the possibility of transient signals when the legs first deployed. The need to guard against such spurious signals was not explicitly included within the the Software Requirements Specification. In consequence, this requirement was not propagated into subsequent test protocols.

Table 10.11 illustrates further benefits of this analysis technique. ECF charts, typically, stretch over many pages. As can be seen from Figure 10.24, this can separate key events during the analysis and testing of a system from the point at which it is presumed to fail. The drafting of XB0114 occurred long before contact was lost with the Polar Lander. ECF charts, such as that shown in Table 10.11, help to trace the impact that distal events and conditions have upon catalytic failures. This is a significant benefit for complex, technological incidents. For example, our analysis of the Polar Lander failure and the associated loss of the Deep Space 2 probes extends to well over fifty nodes. This analysis is still at a relatively high level of abstraction. Several other investigations have produced ECF charts that contain over one thousand events and conditions. In such circumstances, it is essential that analysts have some means of summarising and collating information about the

Event : Premature Shut-down of engines	
What led to the event?	Software did not reset a variable to denote that a spurious touchdown signal had been detected. This variable was read when the touchdown sequence was enabled at forty meters. The lander had an approximate velocity of 13 meters per second, in Martian gravity this accelerates to 22 meters per second at impact.
What went wrong?	Data from the engineering development deployment tests, flight unit deployment tests and Mars 2001 deployment tests showed a spurious reading in the Hall Effect touchdown sensor during landing leg deployment. These spurious signals can continue long enough to be detected as valid. Software that was intended to protect against this did not achieve the intended result. Spurious signals were retained until the sensing logic was enabled at 40 meters from the surface.
How did the barriers fail?	Requirements document (XB0114) did not explicitly state possible failure modes. Software engineers were not told about the transient failures. The system level requirements included a clause that might have alerted engineers to this problem but it was not included in Software Requirements Specification. The transient protection requirement was not, therefore, tested in either the unit or system level tests nor was it looked for in software walk-throughs. There was also an attempt to load balance on the processor so sampling started well before the 40 meter threshold. Product Integrity Engineer for Hall Effect sensors was not present at walk-throughs.
Who was involved in the event?	Software engineers, Product Integrity Engineers.
Is the event linked to a more general deficiency?	Problems in the Polar Lander software for the communications up-link. Software problems also affected Climate Orbiter and Stardust.

Table 10.11: ECF Analysis of the Polar Lander Failure.

key events that contribute to an incident.

Previous paragraphs have used causal factor analysis to drive a more detailed consideration of the events that immediately precede the loss of the Polar Lander and the Deep Space 2 mission. If there was sufficient funding, then investigators would continue the analysis for each event on every path to the incident. If the incident would not have occurred without this event then the supplementary questions in Tables 10.10 and 10.11 would be posed. This approach might be seen to impose unwarranted burdens upon an investigation team. As we have seen, however, it can help to identify new insights into the events leading to high-criticality failures even if other forms of analysis have already been applied. Brevity prevents an exhaustive exposition of this approach. In contrast, Figure 10.25, therefore, presents an ECF chart for the loss of the Climate Orbiter. As can be seen, this diagram integrates the events and conditions from several previous diagrams. These earlier figures included continuation symbols. Figure 10.25 uses these to piece together a more complete view of the incident. As before, however, it is not possible to provide a single legible diagram of all of the events and conditions that were identified by the previous use of change and barrier analysis.

One of the reasons for focusing on Figure 10.25, rather than repeating the causal factor analysis of Deep Space 2 or the Polar Lander, is that it can be used to illustrate the distinction between contextual and causal factors. As before, the analysis starts from the event that precedes the incident. In this case, we must consider whether the incident would still have occurred if the Last signal from MCO (09:04:52, 23/9/99) had not occurred. It seems clear that the incident might still have occurred even if this event had not taken place. If we had omitted this event from our model, we would still have had a coherent explanation of the failure. It, therefore, represents a contextual rather than a causal factor. It is an event that helps our understanding of the incident but it is not necessary to our view of the incident. The analysis, therefore, moves to the event that immediately precedes the previous focus for the analysis. In this case, we must consider whether the incident would have occurred if the Mars Orbital Insertion had not taken place. Again, this event can be omitted without jeopardising the account of the failure. Similarly, the end of the cruise phase is not necessary to a causal explanation of the loss of the Climate Orbiter. The analysis, therefore, moves to the event labelled TCM-5 is discussed but not executed (16-23/9/99).

This event illustrates the complexity of counterfactual reasoning if investigators are not careful about the phrases that are used to label the nodes in an ECF chart. They must determine if the incident would have occurred if it was not the case that TCM-5 is discussed but not executed. The complexity in answering this question stems in part from a mistake in the construction of the ECF chart. As mentioned previously, events should be atomic statements. The previous label refers to both the discussion of the maneuver and to the decision not to implement it. In consequence, Figure 10.25 can be simplified by re-writing this event as It is decided not to execute TCM-5. The discussions surrounding this decision could be shown as an additional, secondary chain of events. It would have been easy to write this chapter with the 'correct' version from the start. This was not done because it is important to emphasise that the development of an ECF chart is an iterative process. It does not guarantee the construction of an 'error free' diagram. In consequence, causal factor analysis provides important checks and balances that can be used to support any causal investigation.

The counterfactual question based on the re-writing of the event now becomes would the incident would have occurred if it was not the case that it was decided not to execute TCM-5? This is equivalent to would the incident would have occurred if it was decided to execute TCM-5? Using the counterfactual question as a test, this event can be considered to have contributed to the failure. The incident need not have occurred if TCM-5 had been executed. A number of caveats can be raised to this argument. For instance, this assumes that that TCM-5 would have been performed correctly. It also assumes that the decision would have been taken when it was still possible to correct the trajectory of the Climate Orbiter prior to insertion. There are further complexities. If we ask the subsidiary question would the ECF chart still represent a plausible path to the incident without the event then it can be argued that the omission of TCM-5 did not cause the incident. It provided a hypothetical means of getting the system back into a safe state. It is, therefore, qualitatively different from the active failures that are addressed in previous paragraphs.

The previous paragraph has argued that TCM-5 is a causal event according to the strict appli-

Event : Ground Based Software uses imperial and not metric units for thruster to compile AMD data file	
What led to the event?	The project Software Interface Specification was not followed nor was their sufficient oversight to detect the incorrect representation of thruster performance.
What went wrong?	Thruster performance data was encoded in Imperial units in the ground based SmallForces routine. This was used to calculate the values that were stored in the AMD_File. Trajectory modellers within the navigation team used this data. They expected it to be in Metric units. As a result, their calculation of the velocity change from AMD events was out by a factor of 4.45 (1 pound of force = 4.45 Newtons) [571]. Key members of the small forces software team were inexperienced. They needed more training on the ground software development process in general and about the importance of the Software Interface Specification in particular. Inadequate training about end-to-end testing of small forces ground software. Failure to identify that the small forces ground software was potentially 'mission critical'.
How did the barriers fail?	SIS not used to direct testing of the ground software. Unclear if this software underwent independent verification and validation. Management oversight was stretched during transition from development to operations and so insufficient attention was paid to navigation and software validation issues. File format problems with the ground software AMD files prevented engineers from identifying the potential problem. Lack of tracking data.
Who was involved in the event?	Ground software development team, Project management, Mission assurance manager (not appointed).
Is the event linked to a more general deficiency?	Software problems affect Polar Lander. Many of these relate to development documents.

Table 10.12: ECF Analysis of the Climate Orbiter Failure.

cation of our counterfactual argument. We have, however, also identified counter arguments. The omission of TCM-5 was not a causal event because even if the decision had been taken to perform this operation there is no guarantee that it would have prevented the incident from occurring. This ambiguity stems from the difficulty of counterfactual reasoning about contingent futures. Not only do we have to imagine there was a decision to implement TCM-5 but we also have to be sure that it would have avoided the incident. The complexity of such arguments has led a number of research teams to apply mathematical models of causation to support informal reasoning in accident investigation [470, 118]. These models attempt to provide unambiguous definitions of what does and what does not constitute a causal relation. They are, typically, based on a notion of distance between what actually happened and what might have happened under counterfactual arguments. A scenario in which TCM-5 was performed and did avoid the incident might be argued to be too far away from the evidence that we have about the actual incident. Such approaches offer considerable benefits; they can be used to prove that different investigators exploit a consistent approach to incident analysis. Unfortunately, the underlying formalisms tend to be unwieldy and error-prone especially for individuals who lack the appropriate mathematical training. A related point is that mathematical definitions of causation are frequently attacked because they fail to capture the richness of natural language accounts. This richness enables investigators argue about whether or not particular events, such as the omission of TCM-5, are actually causal. There would be no such discussion if everyone accepted the same precise mathematical definition! The key point here is that there must be some form of consistency in determining whether or not to explore particular events during any causal analysis. This can either be done by developing strict mathematical rules that can be applied to formal models of causation. Alternatively, they can be drafted as heuristics that can guide less formal analysis by teams of incident investigators. Different forms of ECF tables might be developed to identify any factors that are particularly important for errors of omission [365]. A further alternative might be to ensure that omitted barriers do not appear in the primary event line of an ECF chart because they are explicitly represented by questions in the causal factor analysis. Unfortunately, the documentation associated with existing applications of the ECF approach does not provide any guidance on how this approach might be developed. Instead, there is an emphasis upon the subjective importance of any analysis. There has been no research to determine whether this results in significant inconsistencies between the analysis of different teams of investigators applying the same technique.

Table 10.12 presents the results from applying ECF analysis to **Ground-based software uses imperial not metric units for thruster to compile AMD data file**. This event occurred each time an AMD maneuver altered the Climate Orbiter's trajectory. As can be seen, the use of Imperial units stemmed from a failure to follow the Software Interface Specification. This document required the use of metric units but the development staff received insufficient training to appreciate the significance of this document. As with the previous examples of causal factor analysis, this example also shows how the tables can be used to collate information about an event that might otherwise be distributed throughout an ECF chart. In this case, the Software Interface Specification was not used to guide test case generation. This provides an example of the way in which omitted barriers can be represented within the products of a causal factor analysis, rather than being explicitly introduced into an ECF chart as was the case with the decision not to perform TCM-5.

As before, Table 10.12 identifies some of the individuals and groups who were involved in this event. It also refers to a 'mission assurance manager'. This role had existed in previous missions but no-one performed this role during the Climate Orbiter mission. This illustrates how ECF tables can go beyond the omission of barrier events to also represent the lack of key staff who might have prevented the incident. Finally, Table 10.12 identifies some of the features that are shared between a number of similar incidents. In particular, it refers to the role of development documentation in both the Polar Lander and Climate Orbiter case studies. In the former case, requirements document XB0114 failed to provide programmers with enough information about potential failure modes for the Hall Effect sensors. In the later case, software developers failed to follow the Software Interface Specification because they failed to understand the importance wither of this document or the code that they were writing.

Event	Contextual/ Causal	Justification
Mishap investigation board is established	Contextual	Post-incident event.
Both DS2 probes suffer electrical failure at impact	Causal	The incident would not have happened if this had been avoided.
Forces at impact compromise aft body battery assembly	Causal	The incident would not have happened if this had been avoided. Providing that the RF components were not compromised.
Forces at impact compromise RF components	Causal	The incident would not have happened if this had been avoided. Providing that the battery body assembly was not compromised.
Both DS2 probes impact with the surface	Contextual	Normal or intended behaviour.
Both DS2 probes separate correctly from the MPL	Contextual	Normal or intended behaviour.

Table 10.13: Summary of the ECF Analysis of the Deep Space 2 Incident.

10.3.2 Cause and Contextual Summaries

Causal factor analysis proceeds in the fashion described in previous paragraphs. Investigators iteratively pose counterfactual questions to determine whether each event in an ECF chart can be considered to be causal or not. Table 10.13 summarises the results of this analysis for the loss of the Deep Space 2 probes. As can be seen, there are three causal events: Both DS2 probes suffer electrical failure at impact; Forces at impact compromise aft body battery assembly and Forces at impact compromise RF components. An electrical failure jeopardises the mission if either the aft body battery assembly is compromised or the RF components fail at impact. Each of these events is an element of what Mackie calls a ‘causal complex’ [508]. It is the conjunction of singular causes within the causal complex that leads to a particular outcome. Crucially, the causal complex is sufficient for the result to occur but it is not necessary. There can be other causal complexes. If any of the necessary causal factors within a causal complex are not present then the incident would not have occurred in the manner described.

Table 10.14 extends the previous analysis of the Deep Space 2 probes to account for the loss of the Polar Lander. This identifies three causal factors. Two are relatively straightforward. This incident would clearly have been avoided if the Hall Effect sensors had not generated transient signals. Similarly, the failure would not have happened if the Lander’s engines had not been prematurely cut at 40 meters above the surface. The third event is less easy to assess because it describes the failure of a potential barrier. The software provided some protection against transient signals by rejecting spurious readings from individual sensors. However, it failed to reset the touchdown variable that was used to determine whether the engines should be cut. Table 10.14 argues that this is a causal failure because had the code been written correctly then the incident would not have occurred. This event again illustrates the iterative nature of causal factor analysis.

Even at this advanced stage, it is possible to identify potential improvements to the underlying ECF charts. For example, the analysis presented in Table 10.14 depends on a number of complex counterfactual arguments. These can be simplified by restructuring the underlying ECF charts. For example, the event labelled Software marks individual legs as failed if they show spurious signals but does not reset touchdown indicator at 40 meters (entry +5:16) can be divided into two component events.

Event	Contextual/ Causal	Justification
Mishap investigation board is established	Contextual	Post-incident event.
Premature Shut-Down of engines (40 meters above surface)	Causal	The incident would not have happened if this had been avoided.
Software marks individual legs as failed if they show spurious signals but does not reset touchdown indicator at 40 meters (entry +5:16)	Causal (Barrier)	The incident would not have happened if this had been avoided. This represents a failed barrier because the software does check for spurious signals in individual legs but does not reset the Touchdown indicator.
Radar detects surface of Mars is 40 meters away (entry +5:15)	Contextual	Normal or intended behaviour.
Software marks a touchdown indicator as true if two spurious signals received from the same leg (10-20 milliseconds after deployment)	Contextual	The incident would not have happened if this had been avoided. The software could have disregarded sensor values until some period after leg deployment.
Transient signals possible from Hall Effect magnets when legs first deploy at 1,500 meters (Entry +4:13)	Causal	The incident would not have happened if this had been avoided.

Table 10.14: Summary of ECF Analysis for Polar Lander Incident (Part 1).

One might represent the successful operation of the software defence Software marks individual legs as failed if they show spurious signals. The second event might denote the potential failure Software does not reset touchdown indicator before 40 meters. The former is a contextual event that represents normal or intended behaviour. The latter event can be seen as a causal factor. It represents a failed barrier that might have prevented the incident from occurring had it been correctly implemented.

Table 10.14 summarises the causal and contextual factors that contributed to the loss of the Polar Lander. In particular, it focussed on the potential software failure and its consequent effect of prematurely shutting down the engines while the craft was still some forty meters above the planet surface. Table 10.15 extends this analysis by assessing the events that were used to denote the development and validation of the Lander in previous ECF charts. Two causal events can be identified in this summary: Preliminary design review passed and Launch approved. This analysis again illustrates the practical complexity of counterfactual reasoning about complex failures. For example, it can be argued that both of these events are anticipated within the normal development process and hence should be regarded as contextual rather than causal. The events themselves do not lead to the incident. It is the conjunction of the event together with critical conditions, such as the absence of a system level hazard analysis, that creates a potential failure. Other so-called ‘normal’ events, such as the end of the cruise phase, are not directly associated with such conditions and hence are not considered to be causal. From this it follows that investigators must not only consider the nature of individual events but also the conditions that affect or modify those events in order to determine whether or not they contributed to the causes of an incident.

Tables 10.16 and 10.17 turn from an analysis of the Polar Lander to examine the ECF charts for

Event	Contextual/ Causal	Justification
Last signal from MPL/DS2 (12:02, 3/12/99)	Contextual	Normal or intended behaviour.
Final Trajectory Correction Maneuver (TCM5) begins (05:30, 3/12/99)	Contextual	Normal or intended behaviour.
Cruise phase ends (3/12/99)	Contextual	Normal or intended behaviour.
MPL and DS2 launched (3/1/99)	Contextual	Normal or intended behaviour.
Launch approved	Causal	The incident would not have happened if this had not happened. This could be considered as a normal or intended behaviour. However, the launch should not have been approved without further systems-level analysis and tests.
Development completed	Contextual	Normal or intended behaviour.
Preliminary Design Review passed	Causal	This might be considered a normal or intended behaviour and hence should be contextual rather than causal. However, passing the PDR without further risk management was a causal factor.
Decision to use pulse-mode control	Contextual	This event contributed to the incident because it added to the complexity of the development process and thereby consumed additional design resources.
Decision to use off-the-shelf engines in 4x3 configuration	Contextual	This event contributed to the incident because it added to the complexity of the development process and thereby consumed additional design resources.

Table 10.15: Summary of ECF Analysis for Polar Lander Incident (Part 2).

Event	Contextual/ Causal	Justification
MCO Mishap Investigation Board is formed (15/10/99)	Contextual	Post-incident event.
Operations navigation team consult with spacecraft engineers to discuss discrepancies in velocity change model (27/9/99)	Contextual	Post-incident event.
Last signal from MCO (09:04:52, 23/9/99)	Contextual	Normal or intended behaviour. The signal was lost as the craft passed behind the planet during orbital insertion.
Mars Orbital Insertion begins (09:00:46, 23/9/99)	Contextual	Normal or intended behaviour.
Cruise phase ends (23/9/99)	Contextual	Normal or intended behaviour.
TCM-5 is discussed but not executed (16-23/9/99)	Causal (Barrier)	The failure of a barrier causes problems for counterfactual reasoning because it relies upon subjunctive arguments that may, or may not be justified. In this case, we consider it likely that TCM-5 would have avoided the incident had it been performed.
(File format) anomaly is not reported through Incidents, Surprises, Anomaly system	Contextual (Barrier)	This also depends on a subjunctive argument about whether or not the ISA system might have prevented the incident had it been used. In this case, it is considered that the incident might still have occurred even if the file format anomaly had been reported.
It is apparent that AMD file data is anomalous (N + 7/4/99)	Contextual	Not causal because it created an opportunity to avoid the incident.
File format problems for AMD data is corrected (N/4/99)	Contextual	Not causal because it created an opportunity to avoid the incident.

Table 10.16: Summary of the ECF Analysis of the Climate Orbiter Incident.

the loss of the Climate Orbiter. Table 10.16 identifies a single cause in the events immediately before Mars Orbital Insertion. This relates to the decision not to perform TCM-5. Previous paragraphs have explained how this event can be viewed as causal, if one accepts that TCM-5 is likely to have avoided the incident, or as contextual, if investigators determine that TCM-5 need not have affected the loss of the mission. This illustrates the complexity of informal, subjunctive, counterfactual reasoning. Particular conclusions often depend on the investigators' confidence in a process or device, such as the TCM-5 maneuver. In consequence, the value of structures such as Table 10.16 is not that they simply this difficult form of reasoning. It is, however, that they provide a means of explicitly recording the outcome of such analysis. They also, very importantly, provide a summary justification for any decision to classify an event as either contextual or causal.

Table 10.17 identifies seven causal factors, of which three relate to the failure of potential barriers. The incident would not have occurred if the `SM_Forces` routines had not used Imperial, rather than Metric, units to calculate the values in the AMD file. These values would not have been so critical if engineers had not rejected to use the barbecue mode or if a symmetrical design had been chosen. The failed barriers relate to the lack of independent verification and validation for the `SM_Forces` software. They also stem from the limited number of personnel who made the transition between development and operations. The lack of any a priori hazard analysis early in the development project also removed further protection. The identification of these failed barriers as potential causes again depends upon complex forms of counterfactual reasoning. For example, the small number of development staff being moved into operational roles can only be considered a causal factor if investigators believe that a greater number of development staff would have avoided the problems that affected the mission. It is possible to develop formal models that codify and, therefore, simply counterfactual reasoning. However, these approaches ultimately depend upon investigators determining whether or not such changes in the course of events might have avoided the ultimate failure. The complexity of counterfactual reasoning is, therefore, only partly due to the difficulty of constructing valid arguments. It also stems from the inherent difficulty in constructing arguments that are based on limited knowledge about events that we know did not actually take place.

The previous analysis has a number of important limitations. In particular, it follows the recommended ECF practice of focusing the analysis on events [208, 210]. This creates problems because conditions often provide a common link between many different causal events. Such relationships can be represented in an ECF chart. They can, however, become obscured by the tabular form of analysis that is used to summarise the results of any counterfactual analysis. A further concern is that different investigators may make very different choices when deciding whether or not to represent particular factors as events or conditions. For example, we could introduce a condition which states that requirements document XB0114 does not explicitly consider the failure modes for the Hall Effect sensors. The same omission can also be represented by a number of putative events; Requirements document XB0114 published without failure modes or Decision to omit failure modes from XB0114. These concerns are compounded by the observation that managerial failures are often represented as conditions while individual instances of human error often reveal themselves as discrete events.

A number of approaches can be used to counter-balance this bias towards events. For instance, it is possible to repeat the previous analysis but instead focus upon conditions rather than events. An example of the counterfactual question would then be 'would the incident have occurred if it was not the case that the Climate Orbiter's ground software development staff had limited training in this application domain?'. This approach offers a number of benefits. In particular, it ensures that investigators revisit the many different conditions that can emerge during the previous stages of analysis. This process of cross-checking can help to reveal instances in which the same conditions effect many different aspects of an incident. This approach can, however, also introduce a number of practical difficulties. Almost all of the counterfactual questions that can be applied to the conditions in an ECF chart follow the subjunctive forms that have frustrated our previous analysis of failed barriers. It is very difficult to derive an objective answer to the previous example. How can we determine whether improved training would have avoided the incident? An alternative approach is to use Causal-Context summaries as a form of index into the underlying ECF charts. These diagrams retain the broader conditions that help to shape the context for any incident. In contrast, the summary tables strip out much this detail to focus on the elements of Mackie's causal complexes.

Event	Contextual/ Causal	Justification
Ground-based software uses Imperial and not metric units for thruster to compile AMD data file	Causal	The incident would not have happened if this had been avoided.
Limited independent testing of the ground-based SM_Forces routines	Causal (Barrier)	It is considered likely that the incident would not have occurred if there had been greater independent testing of these routines.
SM_Forces routines are written using imperial and not metric units for thruster performance	Causal	The incident would not have happened if this had been avoided.
Angular Momentum Desaturation events	Contextual	Normal or intended behaviour given the MCO's asymmetric design and the decision to reject the barbecue maneuver.
Systems engineering decision to reject daily 180 degree flip to cancel angular momentum build-up.	Causal	The incident might not have happened if the engineers had decided to perform the 'barbecue' maneuver. However, there remains a degree of doubt that this further navigation problems might have been introduced or gone undetected.
Systems engineering decision to use a solar array that is asymmetrical to the MCO body	Causal	The incident might not have happened if a symmetrical design had been introduced similar to the Global Surveyor.
MCO launch (11/12/98)	Contextual	Normal or intended behaviour.
Minimal number of development staff transition to operations (11-12/98)	Causal (Barrier)	The incident might not have happened if more staff had moved from development to operations.
Decision not to perform an a priori analysis of what could go wrong on the MCO.	Causal (Barrier)	The incident might not have happened if more thought had been given to the problems involved in using the MCO design to achieve the navigation accuracy required by the mission.

Table 10.17: Summary ECF Analysis for Climate Orbiter Incident (Part 2).

Cause-context summary tables and ECF charts together provide a stepping stone towards any subsequent root cause analysis. The following paragraphs address a number of the key issues that must be addressed by any root cause analysis technique.

When to begin? Previous chapters have also argued that the early stages of an investigation are often guided by investigators' working hypotheses about the causes of an incident. It is important, however, that these informal ideas should be explicitly represented relatively early if finite investigatory resources are to be maximised. This requirement must be balanced against the dangers of biasing an investigation towards certain causes. Root cause analysis uses the results of the previous techniques to identify common factors behind causal events. As noted in the previous paragraphs, these common factors may already have been identified as conditions within an ECF chart. It is important to stress, however, that root cause analysis "is not an exact science" [208]. The processes of analysis and investigation often uncover potential root causes that were not considered during previous stages of analysis. It is important, therefore, not to freeze the ECF chart or the cause-context tables during the early stages of any analysis.

How do we validate the analysis? We have argued that ECF charts and cause-context diagrams are 'living' documents that must be updated as new information becomes available. It is important, however, that investigators validate the products of any causal analysis. Typically, this is done through regular, minuted team meetings. Increasingly these are used to approve the publication of draft analysis documents via organisational intranets. They provide shared resources that help to guide the continuing investigation. Such publication and distribution mechanisms help to coordinate investigators' activities but must be protected from public disclosure. Ultimately, the products of any root cause analysis must be approved by the members of an investigation team before a final report can be written. This mechanisms for achieving this agreement depend on the scale of the incident reporting system. In local applications, there may only be a single individual who is available to perform the analysis and draft the report. In larger systems, however, there may be formalised procedures for 'signing off' the products of any root cause analysis. These procedures can involve higher levels of management. This raises serious practical and ethical issues if this final stage of approval is seen as a means of potentially filtering the results of any analysis. Some organisations have guarded against this by allowing senior management only to annotate root cause analyses. They are prevented from altering what has already been written. While this approach offers some protection against undue influence, it does not guard against the myriad of informal pressures that can be brought to bare on an investigation team.

How many root causes? The Department of Energy guidelines state that investigators should identify at least one but probably not more than three or four root causes [208]. This guideline seems to be derived from the pragmatics of incident investigation within particular industries. They do not, however, provide any justification for their suggestion. This is unfortunate. Such a pragmatic limit can be seen as a barrier to organisational learning from any mishap in which there were more than four root causes. Such concerns are exacerbated by the observation that there are often many different ways for an incident to occur. In consequence, there any incident investigation may yield a number of root causes for each of these different scenarios. For instance, the Polar Lander could have been lost because of the premature shut-down of the engines. It might also have been caused by a failure in the separation of the Deep Space 2 probes and the Lander from the cruise stage. It could have been caused by a landing on unfavourable terrain. It might also have been caused by failure in the communications up-link and so on. Each of these scenarios was considered to be plausible by the NASA investigation team. Although each hypotheses yielded a small number of root causes, the cumulative effect of considering many different failure scenarios helped the investigators to identify a significant number of lessons for future missions. This would not have been possible had they stopped at the four or five root causes recommended above. It seems more profitable to view resource constraints as the limiting factor. The extent of any root cause analysis provides a good indication of the perceived criticality of any potential failure.

What are the parameters of the analysis? The ECF guidelines argue that "the intent of the analysis is to identify and address only the root causes that can be controlled within the system being investigated, excluding events and conditions that cannot be reasonable anticipated and controlled, such as natural disasters" [208]. It is clearly difficult to control natural disasters, however, this wide

ranging approach does pose a number of important questions. Previous sections have explained how many local incident reporting systems ‘target the doable’. This can prevent effective action from being taken to address common problems that might affect a number of different local groups. In particular, managerial and organisation constraints may be viewed as outside the control of operational departments. It is, therefore, important that any root cause analysis technique should provide explicit means of addressing these higher-level causes of failure.

The previous paragraphs have described some general attributes of the root cause analysis. They have not, however, provided any guidance about the methods and techniques that might be applied to identify these factors from the mass of information that can be derived from the previous stages of analysis. The following sections, therefore, present two different techniques that can be used to identify root causes from the events and conditions that are described in ECF charts and cause-context tables.

10.3.3 Tier Analysis

Tier diagramming is a root cause analysis technique that focuses on those levels of management that have the responsibility to correct potential problems. It is one of several techniques, including Pate-Cornell’s ‘Accident Analysis Framework’, that explicitly force investigators to consider organisational factors as the initial root causes of many failures [667]. Each row in one of these diagrams refers to a different level of management within an organisation. They are intended to represent levels of organisational responsibility that range from the operator up to senior management. The columns in a tier diagram list the causal factors that are derived from the Causal factor analysis together with any higher-level root causes that may or may not be identified. This is illustrated by Table 10.18. It is important to note, however, that this a generic template that must be tailored to reflect the organisations that are involved in a particular incident. Each causal factor is assigned to a tier of management responsibility. This is intended to help identify any common links between causal factors that relate to particular levels in an organisation. For instance, a failure in supervision would be exposed by a number of causal factors that cluster around this level in the tier diagram. This is intended to offer a number of benefits to any incident investigation. In particular, it helps to focus any root cause analysis on the deeper organisational causes of failure [702]. The tabular format also helps to structure an investigation around concepts, or groups, that have a clear organisational meaning for those involved in an incident. This is important because many incident reports often talk in vague terms about a ‘failure in safety culture’ without grounding these observations in the activities of particular organisations and groups. A further benefit is that responsibility is explicitly assigned for each root cause and causal factor. These judgements provide a focus for subsequent discussion and can, ultimately, help to form the recommendations for future practice.

Tier	Causal Factors	Root Cause
5: Senior Management		
4: Middle Management		
3: Lower Management		
2: Supervision		
1: Workers Actions		
0: Direct Cause		

Table 10.18: Format for a Tier Diagram [208].

Different tier diagrams are drawn up for each of the organisations that is involved in an incident. In our case studies, therefore, we would anticipate separate tier diagrams for NASA Headquarters and for NASA JPL and for the subcontractor LMA. It is also possible to refine such diagrams to look at different groups and teams within each organisation. For instance, it is possible to distinguish management tiers within the development process of the Climate Orbiter from operation groups. Tier diagramming, typically, begins with the organisation that is most closely involved in the incident.

The first diagram in both the Polar lander and Climate Orbiter case studies would focus on the LMA operational teams. Further diagrams would then represent the contractor organisation for which LMA was subcontracting. In particular, tier diagrams should also represent any organisations that are involved in the oversight or regulation of the contractor's and subcontractor's activities. Tier diagramming, therefore, has two prerequisites. Firstly, investigators must have already identified a number of potential causal factors using techniques such as causal factor analysis. Secondly, they must also have a clear understanding of the management structures that characterise the organisations involved in an incident. Once this information is available, the analysis proceeds in the following stages:

1. Develop the tier diagram. Create a tier diagram that reflects the management structure of the organisation being considered.
2. Identify direct causes. Examine the cause-context summaries to identify any catalytic events that cannot be directly associated with operators or management activities. Enter these along the direct cause row, shown in Table 10.18. Repeat this process for any conditions that are associated with these causal events in an ECF chart. Initially, this tier might contain events that describe the failure of process components or problems due to the contamination of raw materials. As analysis progresses, however, it is likely that most of these direct causes will be associated with other tiers in the diagram. For instance, component failures may be due to a managerial failure to ensure an adequate maintenance regime. Similarly, the contamination of raw materials can be associated with acquisitions and screening policies.
3. Identify worker actions. For each causal factor in the cause-context summary, ask whether or not they stemmed directly from 'worker actions'. A number of guidelines can be proposed to direct this stage of the analysis. For instance, the US Department of Energy has developed a number of questions that are intended to help determine whether or not a causal factor should be associated with worker actions [208]. These include whether or not the worker's knowledge, skills and abilities were adequate to perform the job safely. They also ask whether the worker understood the work that was to be performed. As with direct causes, these actions often raise questions about the performance of other groups in a tier diagram. The worker's lack of understanding may be due to an inadequate training regime. Investigators must, therefore, ask whether or not the worker was solely responsible for the causal factor. If the answer is no then investigators must move the event to a higher tier in the diagram. As before, investigators must also introduce any associated conditions into a tier diagram if they provide necessary additional information about causal events.
4. Analyse remaining tiers. The analysis progresses in a similar fashion for each tier. The intention is to place each causal factor as high up the diagram as possible. Ultimately, as we have seen, all incidents can be associated with regulatory problems or a failure in oversight. It is important, however, to balance this observation about ultimate responsibility against the need to identify those levels in an organisation that are most directly responsible for certain causal factors. As mentioned in the previous paragraph, this is most often done by developing analytical guidelines. These guidelines help investigators to assess whether or not a causal factor can be associated with a particular tier in the diagram. They are, in turn, typically derived from the safety cases that justify the operation of an application process. For instance, if middle management has an identified responsibility to ensure the operation of an incident reporting system then it is possible to place any causal factor that relates to the failure of such a system at this level in a tier diagram.
5. Identify links. After all of the causal factors and associated conditions have been entered into a tier diagram, investigators can begin to look for common factors. As with the previous stages in this form of analysis, the success of this activity depends upon the skill and expertise of the investigator. This, in turn, can have a profound impact on the course of any investigation. As Lekberg notes, the previous background and training of an investigator can have a profound impact on the results of their analysis [484]. The key point is not, however, to eliminate these

individual differences but to use the tier diagram as a means of explicitly representing the key stages in any root cause analysis. Other investigators can then inspect these diagrams to identify other connections between causal factors or, if necessary, to argue against proposed links. Investigators can use different colours or symbols to denote those causes that are considered to be linked.

6. Identify root causes. Compare each of the causal factors in the tier diagrams against the definition of a root cause. A root cause is distinguished by Lewis' counterfactual argument that if A and B are states (conditions) or events, then A is a necessary causal factor of B if and only if it is the case that if A had not occurred then B would not have occurred either [490]. This is essentially the same requirement that was used to distinguish causal from contextual factors in the causal factor analysis. They can also be thought of as causal factors that, if corrected, would prevent recurrence of the same or similar incidents. We would also impose an additional requirement based on Mackie's distinction between general and singular causes [508]. Root causes must address a class of deficiencies, rather than single problems or faults. Correcting a root cause not only prevents the same incident from recurring but also solves deeper line management, oversight and management system deficiencies that could cause or contribute to future mishaps [208]. If a causal factor meets these criteria then an additional entry can be made to denote this finding in the third table of the tier diagram, illustrated in Table 10.18. Investigators must, therefore, compose a root cause 'statement' to summarise each of the causal factors groupings that were identified in the previous stage of analysis.

Root cause analysis can reveal events and conditions that were not represented on ECF charts, ECF tables or cause-context summaries. . These must be added to ensure consistency between these various products of a root cause analysis. It should also be noted that one tier diagram may provide input for another. For instance, if the upper management of a contractor was responsible for a particular root cause then the regulator and supervisory organisation may share responsibility for that particular root cause if there is a deficiency in the directives given by those organisations.

The remainder of this section applied the tier diagramming approach to identify root causes for both the Polar Lander and the Climate Orbiter case studies. This analysis begins by identifying the relevant management and organisation structures that were involved in this incident. The Mars Independent Assessment Team have provides information about the internal management structures within NASA headquarters and within JPL [570]. Unfortunately, it can be less easy for investigators to obtain detailed information about subcontractors' management structures even in the aftermath of a serious incident. The subsequent analysis, therefore, must also exploit a number of inferences about the reporting structures that characterised the day to day operation of the Mars Surveyor projects.

Figure 10.26 illustrates the complexity of the management structures that were involved in the Mars Program at NASA Headquarters. Not only do such organisational features complicate any tier analysis, they also had a significant impact on the loss of the Polar Lander and the Climate Orbiter. During the initial formation of the program, the JPL Program Manager had to deal with the Advanced Technology and Mission Studies Division. During implementation, they interacted with the Mission and Payloads Development Division. For the operational phase of the program, the JPL Program Manager dealt with the Research and Program Management Division. During all of this the manager must also interact with the Science Board of Directors. These various channels of communication between NASA headquarters staff and the JPL Mars Program Manager caused problems for both organisations. The independent assessment team found that "ineffective communication between JPL management and NASA Headquarters contributed to an unhealthy interface and significant misunderstandings in conducting the Mars Surveyor Program" [570]. NASA Headquarters believed that they were articulating program objectives, mission requirements and constraints. JPL management interpreted these as non-negotiable demands over costs, schedules and performance requirements. Concern about losing contracts and funding also prevented JPL management from effectively express their concerns to NASA Headquarters about programmatic constraints. The independent assessment team also concluded that NASA Headquarters did not seem receptive to receiving bad news.

JPL's Mars Program Office initiated the Mars 98 project and was responsible for planning, program advocacy and flight project development between 1994 and 1996. The roles and responsibilities of this office were, however, interpreted differently in the JPL Mars Program Office and the NASA Headquarters sponsoring office. This led to several conflicts about mission direction that ultimately diverted management resources away from mission development. These difficulties illustrate an important practical barrier to tier analysis. One of the precursors to an incident may be the breakdown of management structures. The roles and responsibilities of each level of the table may, therefore, be very difficult to distinguish: "individual projects were not developed or managed within a clearly defined overall framework that identified interdependencies and risk management strategies" [570].

In 1996, NASA Headquarters delegated full program management authority to the NASA Centers. JPL, therefore, created a Mars Exploration Directorate that reported directly to the Laboratory Director. This directorate assumed responsibility for program management and assumed most of the duties that have previously been associated with the NASA Headquarters sponsoring office. One consequence of this reorganisation was that JPL's Mars Exploration Directorate lost a single point of contact at Headquarters. In August 1996, the management structure of the Mars programs was further complicated by the announcement that potential signs of life had been found on a meteorite that was assumed to have come from Mars. The heightened public interest led to further changes in JPL's organisation. An increased emphasis was placed on robotic exploration to support the long-term needs of Human Exploration. These missions were managed by a different part of Headquarters

JPL responded to these changes in priorities by partially reorganising its own management structure in 1998. This was followed by wider changes in 1999. JPL amalgamated its space and Earth science teams into a single directorate. The intention was to coordinate the management of an increased number of programs and projects in both of these areas. The Mars Program Manager no longer reported to the Laboratory Director as a separate, independent entity. Project managers were to report at a lower level. Figure 10.27 illustrates the organisational structure of the JPL Space and Earth Sciences Programs Directorate after the 1999 reorganisation. The Mars projects are shown among sixty-eight other projects in the third tier of management. They are, therefore, isolated from the direct reporting structures of senior JPL management. Although Figure 10.27 represents the 1999 reorganisation, the independent assessment team argued that this reflects the project isolation that contributed to the failure of the Mars'98 project.

The previous paragraphs have summarised the management structures within NASA headquarters and within JPL. They have also argued that the dynamism of many organisations can create significant problems when applying tier analysis to real-world management structures. The different teams and individuals who are associated with different levels in a tier diagram may change as organisations attempt to adapt to the pressures that are created by many high-technology projects. One solution would be to develop a number of tier diagrams to represent these different changes in project management. An alternative approach is to exploit a relative abstract classification of organisational structures, similar to those shown in Figure 10.18 and then provide more detailed information to support the interpretation of those categories at particular stages of the incident.

A number of further challenges complicate the development of tier diagrams. In particular, it may not be possible for the investigators from one organisation to gain access to detailed information about the management of another organisation. As we have seen, it is relatively easy to access documentation about NASA management structures. It is far harder to find comparable information about the organisation of the commercial subcontractors. In consequence, investigators may be forced to exploit the more generic tiers that were introduced in Table 10.18. Even if this approach is exploited, investigators face a number of further problems. For example, if there are several organisations involved in an incident then they must determine which causes relate to which tier diagram. This can partly be based on any existing project documentation, however, it also requires considerable skill and judgement on the part of individual investigators. For example, the following quote illustrates how LMA were responsible for the development of the Mars Surveyor programme. JPL staff were involved in some of these activities but they also provided higher level management functions:

"The Mars Surveyor program'98 Development Project used a prime contract vehicle

to support project implementation. Lockheed Martin Astronautics (LMA) of Denver, Colorado was selected as the prime contractor. LMA's contracted development responsibilities were to design and develop both spacecraft, lead flight system integration and test, and support launch operations. JPL retained responsibilities for overall project management, spacecraft and instrument development management, project system engineering, mission design, navigation design, mission operation system development, ground data system development, and mission assurance. The MSP 98 project assigned the responsibility for mission operations systems/ground data systems development to the Mars Surveyor Operations Project, LMA provided support to Mars Surveyor Operations Project for mission operations systems/ground data systems development tasks related to spacecraft test and operations." [565]

This quotation illustrates the practical difficulties that are involved in separating out the responsibility that each organisation might assume for certain causes of safety-critical incidents. In consequence, the following tables represent one particular viewpoint. They act as a focus for subsequent discussion rather than a unique assignment of causal factors to particular management layers in each of the organisations.

Figure 10.19 provides an initial assignment of causes to various layers within the contractor organisation. In addition to these causal factors, identified in the cause-context summaries, it is also possible to introduce conditions that are also perceived to have contributed to the incident. As mentioned, these conditions can represent longer term factors that cannot easily be represented as discrete events and so may be overlooked by the previous forms of analysis. For instance, previous ECF charts identified the way in which some project requirements were not passed on in sufficient detail. This was shown as a condition labelled *Requirements are not passed on in sufficient detail nor are they backed by an adequate validation plan* in Figure 10.11. This created problems because individual project managers had to interpret what was admissible in pursuit of the objectives set by *Faster, Better, Cheaper*. Figure 10.19, therefore, introduces a number of similar conditions into the tier diagram.

It is important to note that Figure 10.19 represents the management structure that was in place at JPL between 1994-1996. It was during this period that JPL's Mars Program Office initiated the Mars 98 project and was responsible for planning, program advocacy and flight project development. As noted in previous sections, tier analysis is complicated by the fact that the management tiers were altered several times during the project lifecycle. Figure 10.27, shown previously, illustrates the JPL management structure that was put in place from 1996. A new Mars Exploration Directorate was created within JPL to coordinate many of the activities that were previously performed by NASA Headquarters and so are not considered in Figure 10.19.

Figure 10.19 illustrates the way in which tier analysis tends to associate root causes with the higher levels of management. This is a natural consequence of the iterative process that is used to analyse each causal factor; the intention is to place each causal factor as high up the diagram as possible. This is an important strength of the technique. The investigators' attention is drawn away from individual instances of operator error. Undue emphasis may, however, be placed on individuals at higher levels within an organisation. This is inappropriate if operational responsibility is devolved to lower levels within the management structure. Under such circumstances, any root cause for the failure might have to be associated with several different levels within an organisation.

The distribution of responsibility within an organisation is illustrated in Figure 10.19 by root causes at both senior and middle management level. Although senior personnel provided insufficient guidance on the implementation of NASA's *Faster, Better, Cheaper* strategy, middle management might still have fought to obtain adequate resources. This also illustrates the subjective nature of tier analysis. It can be argued that these two root causes are so closely linked that they should be amalgamated into a single higher-level description. If Senior Management had provided strong guidance about the implications of the *Faster, Better, Cheaper* strategy for design and validation then Middle Level Management would have had less need to fight for additional resources. On the other hand, it can be argued that these root causes should be distinct because Senior Management must rely on their colleagues to provide adequate information about the operational implication of accepting such tight resource constraints. Similarly, there are some causal factors in Figure 10.19 that

Tier	Causal Factors	Root Cause
Senior Management	<p>Requirements are not passed on insufficient detail nor are they backed by an adequate validation plan.</p> <p>Decision not to perform an a priori analysis of what could go wrong on the MCO.</p> <p>Limited independent testing of the ground-based SM_Forces routines.</p>	<p>No documented guidance on implementing Faster, Better, Cheaper prevented project managers from resisting pressures to cut costs/schedules that might compromise mission success.</p>
Middle Management	<p>Minimal number of development staff transition to operations (11-12/98).</p> <p>SM_Forces routines are written using imperial and not metric units for thruster performance.</p>	<p>Lack of resources for the Mars Surveyor Program limited the number of staff available and may also have prevented those staff from receiving adequate training on critical aspects of the mission.</p>
Lower Management	TCM-5 is discussed but not executed (16-23/9/99)	
Supervision		
Workers Actions	<p>Systems engineering decision to reject daily 180 degree flip to cancel angular momentum build-up.</p> <p>Systems engineering decision to use a solar array that is asymmetrical to the MCO body</p>	
Direct Cause	Ground-based software uses Imperial and not metric units for thruster to compile AMD data file	

Table 10.19: LMA Tier Diagram for the Climate Orbiter Mission.

could have been represented as root causes. The decision not to implement TCM-5 is an example of one such event. If this maneuver had been implemented then the incident could have been avoided. The lack of preparation for this maneuver and the consequent decision not to implement it might, in combination with other factors, lead to future incidents. The key point here is that either approach would represent a valid application of tier analysis. The output of this process depends upon the skill, expertise and viewpoint of the investigator. It, therefore, must be carefully validated by peer review. One means of validating our analysis would be to compare Figure 10.19 with the output of an independent tier analysis performed by another investigator. There may, however, be more general biases that are introduced by the use of this particular form of analysis. An alternative means of validating these findings is to compare the results of our analysis with those obtained by investigators using other approaches. For example, the following section will repeat the analysis of our case studies using Non-compliance classifications. For now it is sufficient to summarise the findings of the Mars Program Independent Assessment Team Report. They used a range of less structured techniques to derive the following conclusions about contractor involvement in the root causes of the incident:

“(NASA, JPL, and LMA) have not documented the policies and procedures that make up their Faster, Better, Cheaper approach; therefore, the process is not repeatable. Rather, project managers have their own and sometimes different interpretations. This can result in missing important steps and keeping lessons learned from others who could benefit from them... Mars 98 had inadequate resources to accomplish the requirements. Through a combination of perceived NASA Headquarters mandates and concern for loss of business, JPL and LMA committed to overly challenging programmatic goals. The JPL management perception was that no cost increase was permissible and the aggressive pricing strategy adopted by LMA exacerbated the problem. The pressure of meeting the cost and schedule goals resulted in an environment of increasing risk in which too many corners were cut in applying proven engineering practices and the checks and balances required for mission success... Inadequate project staffing and application of institutional capability by JPL contributed to reduced mission assurance. Pressure from an already aggressive schedule was increased by LMA not meeting staffing objectives early in the project. This schedule pressure led to inadequate analysis and testing. An additional important role for senior management, whether at NASA, JPL, or LMA, is to ensure the establishment of, and compliance with, policies that will assure mission success. For example, these policies should address design (at the component, system, and mission life cycle level), test and verification, operations, risk management, and independent reviews.” [570]

As can be seen, several of the themes identified by the Mars Program Independent Assessment Team are summarised as root causes in the tier analysis of Figure 10.19. There are some differences. In particular, the team’s report brings together many of the factors that we have identified and links them to the contact management’s perception of project risk. Our analysis was performed prior to reading this document. With this additional insight, however, it would be possible to reformulate the previous diagram to reflect these more general concerns. This again reflects the point that root cause analysis is an iterative process. ECF charts, cause-context summaries, tier analysis are all artifacts that help to document the path towards a causal analysis. They do not replace the skill and expertise of the investigators nor do they ‘automate’ key stages of the analysis.

Figure 10.20 builds on the previous analysis by examining the root causes of the Climate Orbiter failure from the perspective of the JPL management structure. Unlike the contractor organisations, more can be identified from the published documentation about management structures within this organisation. As mentioned previously, JPL retained responsibilities for “overall project management, for spacecraft and instrument development management, for project system engineering, mission design, navigation design, mission operation system development, ground data system development and mission assurance” [571]. From this it follows that JPL staff were ultimately responsible for the development and testing of the navigation software. It can, therefore, be argued that some of the causal factors associated with navigation systems development should be removed from Fig-

ure 10.19 The contractor was not responsible for overseeing this aspect of the mission. These factors have been retained because the NASA investigators commented on the difficulty of making such precise distinctions, staff often could not reply to questions such as ‘who is in charge?’ or ‘who is the mission manager?’ [571].

Figure 10.20 shows how causal factors affect several of the organisations that are involved in any incident. This diagram presents many of the events and conditions that were identified in the tier analysis for LMA staff. However, the supervisory and managerial role of JPL staff is reflected by the way in which many of these causal factors are associated with different levels in the management structure. For instance, the event TCM-5 is discussed by not executed was associated with lower levels of management within the contractor organisation but is associated with the program management in JPL. The Flight Operations Manager should have polled each subsystem lead to ensure that they had reviewed the data and believed that the Climate Orbiter was in the proper configuration for the event. [571] However, this protocol had not been developed nor had any manager been explicitly identified to lead this decision making process. It might, therefore, be argued that responsibility rested with the JPL program manager, as shown in Figure 10.20.

Figure 10.20 also illustrates the manner in which tier analysis can expose different root causes for similar causal factors within different organisations. For example, the inadequate risk analysis and the lack of development staff who transitioned into operations might indicate a degree of complacency on the part of the JPL management team. The NASA investigators found evidence of a perception at JPL that “orbiting Mars is routine” [571]. This perception was based on previous mission successes. However, it resulted in inadequate attention being paid to navigation risk mitigation.

Figure 10.20 also illustrates the way in which tier diagram must account for the relationship between the management structure that is being considered and any other organisations that are involved in an incident. In this case, the insular relationship between JPL and the contract organisation is identified as a root cause behind the lack of independent testing and inadequate risk assessment. This analysis raises a number of structural properties about our use of the tier diagrams in Figure 10.20. As can be seen, causal factors and root causes are associated with different levels of management. No distinction is made between these causes. For instance, only two out of the three causal factors at the top levels of the JPL management structure are associated with the insularity, mentioned above. Similarly, we have not shown how causal factors at various levels in a tier diagram might contribute to a root cause. Additional annotations could be introduced to represent this information. Care must be taken if the resulting diagrams are not to become illegible.

As before, we can compare the results of the tier analysis with the findings of the Mars Program Independent Assessment Team. The root cause analysis illustrated in Figure 10.20 is based on a subset of the evidence that was available to this investigation team. Our analysis was, however, done prior to reading their account:

“The JPL/Lockheed Martin Astronautics interface for Mars 98 was characterised by a positive, close working relationship between the JPL and LMA project managers and their offices. However, this relationship had a negative, insular effect when accepting excessive risk... Inadequate project staffing and application of institutional capability by JPL contributed to reduced mission assurance. Pressure from an already aggressive schedule was increased by LMA not meeting staffing objectives early in the project. This schedule pressure led to inadequate analysis and testing... The team found multiple examples of ineffective risk identification and communication by both JPL and LMA. Compounding this, JPL and LMA each deviated from accepted and well-established engineering and management practices. Risk identification and any significant deviations from acceptable practices must be communicated to the customer in an open, timely, and formal fashion.” [570]

It is difficult in the aftermath of such an incident to be sure that this analysis has not biased my interpretation of the incident. The findings of the Mars Program Independent Assessment Team were publicised in press accounts. They are also referenced in the pages that provided access to on-line versions of primary sources that were used in our analysis. Any comparison between the results of our tier analysis and the assessment team’s report cannot, therefore, be regarded as an independent

Tier	Causal Factors	Root Cause
5: Senior Management (JPL Laboratory Director and Mars Program Office Director)	<p>Minimal number of development staff transition to operations (11-12/98)</p> <p>Limited independent testing of the ground-based SM_Forces routines</p> <p>Decision not to perform an a priori analysis of what could go wrong on the MCO.</p>	<p>Feeling that orbiting Mars in routine.</p> <p>Insular relationship with LMA prevented adequate risk assessment and mitigated against independent reviews.</p>
4: Middle Management (Climate Orbiter Project Manager)	TCM-5 is discussed but not executed (16-23/99)	
3: Lower Management (Flight Operations Manager/Flight Development Manager)	<p>SM_Forces routines are written using imperial and not metric units for thruster performance.</p> <p>Systems engineering decision to reject daily 180 degree flip to cancel angular momentum build-up.</p> <p>Systems engineering decision to use a solar array that is asymmetrical to the MCO body</p>	

Table 10.20: JPL Tier Diagram for the Climate Orbiter Mission.

or formal validation of the root causes analysis. In contrast, Figure 10.20 simply illustrates that it is possible for some of the independent assessment team’s findings to be represented within a tier diagram. It is also important to identify the differences between our ECF/tier analysis and the findings of the independent assessment team. In particular, the root causes in Figure 10.20 do not address the communications problems that existed between JPL and NASA headquarters. The Mars Program’s Independent Assessment Team report emphasised that these problems prevented JPL management from gaining a clear understanding of the resource implications behind the Faster, Better, Cheaper strategy. These concerns are, however, represented in Table 10.21 that presents a tier analysis of NASA headquarter’s involvement in the loss of the Climate Orbiter.

Tier	Causal Factors	Root Cause
5: Senior Management (Board of Directors, Science)	Project oversight problems stem from complex relationship between JPL and LMA (and NASA HQ)	Failure to communicate the mission implications of the Faster, Better, Cheaper strategy.
4c: Middle Management (Associate Administrator, Office of Space Science)		
4b: Middle Management (Science Chief of Staff)	Lack of managerial leadership in promoting responsible attitudes to Incidents, Surprises and Anomaly reporting	Failure to communicate the importance of expressing concerns both about specific implementation issues as well as resource/management problems.
4a: Middle Management (Advanced Studies Division, Mission Development Division, Research and Program Management Division etc)	Requirements are not passed on in sufficient detail nor are they backed by an adequate validation plan	

Table 10.21: NASA HQ Tier Diagram for the Climate Orbiter Mission.

Figure 10.21 illustrates the way in which investigators can use both the conditions and the events in an ECF chart to support any subsequent tier analysis. In this case, NASA headquarters had little direct involvement in the events that led to the loss of the Climate Orbiter. Investigators would, therefore, have considerable difficulties in constructing a root cause analysis that was based solely upon such direct involvement. In contrast, it can be argued that NASA headquarters played an important role in establishing the conditions that led to this incident. Figure 10.21 therefore goes beyond the causal events that were considered in previous tier diagrams to look at the conditions that were identified in early ECF charts of the Climate Orbiter incident, such as Figure 10.8. This example is typical of tier diagrams that consider the role of regulatory or supervisory organisations in such failures. It is also important to note that such factors are often omitted from some reports of an incident. For example, the initial report into the Climate Orbiter contained no reference to the involvement of NASA headquarters at all [565]. This is justified by the initial focus on the direct causes of the incident. The subsequent report into Project Management in NASA by the Mars Climate Orbiter, Mishap Investigation Board only contained four references to NASA headquarters [571]. None of these references described any potential inadequacies that might have led to the incident. In contrast, the Mars Program Independent Assessment Team that was supported by

NASA made approximately fifty references to the role played by headquarters [570].

The findings from the Independent Assessment Team can again be compared with the root causes that have been identified using tier analysis. Such a comparison reflects some of the limitations of this approach when applied to the less direct causes of an incident or accident. The following excerpts summarise the results of the independent enquiry:

“ Through a combination of perceived NASA Headquarters mandates and concern for loss of business, JPL and LMA committed to overly challenging programmatic goals. The JPL management perception was that no cost increase was permissible and the aggressive pricing strategy adopted by LMA exacerbated the problem... NASA Headquarters thought it was articulating program objectives, mission requirements, and constraints. JPL management was hearing these as non-negotiable program mandates (e.g., as dictated launch vehicle, specific costs and schedules, and performance requirements)... The result was that JPL management did not convey an adequate risk assessment to NASA Headquarters. What NASA Headquarters heard was JPL agreeing with and accepting objectives, requirements, and constraints. This communication dynamic prevented open and effective discussion of problems and issues. JPL management did not effectively express their concerns to NASA Headquarters about programmatic constraints, and NASA Headquarters did not seem receptive to receiving bad news... In this case, JPL and NASA Headquarters communications were inadequate, in part because JPL was concerned that Headquarters would perceive JPL concerns about programmatic constraints negatively; JPL did not want to antagonise the customer. NASA Headquarters was rigid in adhering to unrealistic constraints. Communication between JPL and NASA Headquarters was impeded by a cumbersome and poorly defined organisational structure within the Office of Space Science.” [570]

Our use of tier analysis did not reveal many of the causal factors that are identified in the Mars Program Independent Assessment Team’s report. For instance, the previous tables did not identify the communications problems that led JPL to interpret Headquarter’s objectives as non-negotiable program mandates. On the other hand, the tier analysis associated a failure to encourage the use of Incident, Surprises and Anomaly reporting with Headquarters management. A number of different explanations can be proposed for such apparent differences. The first is that the subjective nature of root cause analysis, even when supported by ECF charts and tier analysis, makes it likely that different teams of investigators will focus on different aspects of an incident. It is hardly surprising, given the content of this book, that our analysis should have identified the failure of the reporting system as a root cause! A second potential explanation for these apparent differences is that the results of the tier analysis are strongly influenced by the use of ECF charts during the initial stages of an investigation. This technique encourages analysts to focus on particular events rather than on the organisational factors that create the conditions for an incident. It is important to remember, however, that this initial focus is broadened by barrier and change analysis. Both of these techniques help to ensure that causal factor analysis does look beyond the immediate events that contribute to an incident. A third explanation for the differences between the products of our tier analysis and the organisational analysis of the independent assessment team is that each of these investigations had different objectives. Our intention in identifying the root causes of the Climate Orbiter incident was to demonstrate that tier analysis could be used to identify root causes at different levels of management in each of the organisations that were involved in the incident. In contrast, the Mars Program Independent Assessment Team was more narrowly focussed on the structure and organisation of NASA’s Mars Program. It therefore provides only a cursory examination of the direct events leading to the failure and certainly does not approach the level of detail shown in previous ECF charts.

The previous paragraphs have shown tier analysis can be used to identify root causes amongst the conditions and events that are derived from a causal factor analysis. An important strength of this approach is that it focuses the investigators attention on the higher levels of management within the organisations that are involved in an incident. Tier analysis also helps to explicitly distinguish generic causes, i.e., factors that might result in future failures, from the more specific

causal factors that characterise a particular incident. Previous paragraphs have also identified a number of potential weaknesses. Tier analysis may be unnecessarily restrictive if it relies on causal factor analysis as a means of identifying potential causal factors. Unless this technique is used in conjunction with a broad ranging change or barrier analysis then it can be difficult to identify all of the ways in which organisational factors might contribute to an incident. Tier analysis also relies entirely upon the subjective skill of the investigator. It is possible to annotate tier diagrams in a flexible manner but they must be supported by prose descriptions if other investigators are to understand the detailed justification for identifying particular root causes from a mass of other causal factors. These descriptions are important because without them it will be difficult to validate the output from any tier analysis.

10.3.4 Non-Compliance Analysis

Rather than repeat our application of tier analysis for the Mars Polar Lander incident, this section presents an alternative form of root cause analysis. Non-compliance classification focuses on three different forms of non-compliance. The first relates to situations in which individuals *don't know* that they are violating an accepted rule or procedure. This occurs if workers receive inadequate training or if they are not informed about changes in applicable regulations. The second classification deals with situations in which individuals and teams *can't comply*. This occurs if operators or managers are denied the necessary resources to meet their obligations. The final classification relates to situations in which there is a decision not to follow rules and procedures. Individuals and teams may explicitly or implicitly decide that they *won't comply* with an applicable regulation. Table 10.22 summarises the more detailed categories that investigators must consider for each of these possible situations [208].

Don't Know:	
Never Knew	Poor training or a failure to disseminate regulations to the appropriate recipients.
Forgot	Individual factors, inadequate reminders or unrealistic assumptions on the part of an organisation about what can be recalled, especially under stress.
Didn't understand	Lack of experience or of guidance in how to apply information that has already been provided.
Can't Comply:	
Scarce Resources	Often used to excuse non-compliance. Investigators must be certain that adequate resources were requested.
Impossible	Organisations may impose contradictory constraints so that it is impossible to satisfy one regulation without breaking another.
Won't Comply:	
No penalty or no reward	There may be no incentive to comply with a requirement and hence there may be a tendency to ignore it.
Disagree	Individuals and groups may not recognise the importance of a requirement and so may refuse to satisfy it. Local knowledge may suggest that a regulation threatens safety.

Table 10.22: Root Cause Taxonomy within Non-Compliance Analysis.

The US Department of Energy recommends non-compliance analysis as a means of extracting root causes from the mass of more general causal factors that are derived from causal factor analysis [208]. The causal events that are identified using the counterfactual analysis of previous sections are associated with one of the categories shown in Table 10.22. It is worth recalling that causal factors are distinguished using the counterfactual question; would the incident have occurred if this event or condition had not held? Root causes satisfy the additional condition that they must represent a more general cause of future failures. Non-compliance analysis can be used to distinguish root causes from causal factors because each of the categories in Table 10.22 corresponds to a pre-defined set of more general root causes. By classifying a causal factor according to one of these categories, investigators are encouraged to recognise the wider problems that may stem from the associated root causes. Causal factors that fall into the *don't know* class represents a failure in the training and selection of employees. The *can't comply* class represents root causes that stem from resource allocation issues. Causal factors associated with the *won't comply* class represents a managerial failure to communicate safety objectives. For example, previous sections have used causal factor analysis to identify a number of causal factors that may have contributed to the loss of the Climate Orbiter. These included the observation that Ground-based software uses Imperial and not Metric units for thruster performance during calculation of the AMD data file. The programmers failed to follow the recommended practices that were outlined in the Software Interface Specification. Non-compliance analysis might, therefore, conclude that the software engineers never knew about this document, that they did know about it but forgot to use it or that they did not understand its relevance to the development of mission critical software. These classifications all refer to an underlying root cause; employees were not adequately trained to recognise the importance of such documents. In consequence, any remedial actions should not focus simply on the Software Interface Specification but on the more general need to ensure that software engineers have an adequate understanding of the development practices that are outlined in this and similar documents.

This approach offers a number of potential benefits for organisations whose activities are governed by well-documented guidelines, standards and regulations. Some of these documents even provide investigators with advice about how to detect the symptoms of non-compliance. For example, JPL produced a series of documents on NASA recommended practices that explicitly state what might happen if projects fail to follow the guidelines:

“Impact of Non-Practice: The performance of the delivered product may be compromised if the hardware imposed limitations are not evaluated early in the design phase. Once the hardware is delivered, it is too late to select an alternative radio architecture, and there are few opportunities to mitigate the impact of any constraints on radio performance. Lacking insight into RF hardware characteristics, test engineers may waste valuable engineering hours determining the basis for the variance between expected and observed performance. For flight projects, costly problem/failure reports and project waivers will likely be processed due to the lack of an early understanding of hardware limitations.” [579]

There are, however, a number of practical problems that complicate the use of non-compliance analysis as a means of identifying more general root causes from the causal factors that are identified during a causal factor analysis. Firstly, the more general root causes that are associated with the categories in Table 10.22 cannot hope to cover all of the potential root causes of adverse incidents in many different industries. In contrast, this form of analysis directs the investigators' attention towards a very limited set of factors associated with training, with resource allocation and with the communication of safety priorities. This direction can either be seen as a useful heuristic that helps to ensure consistency between analysts or as a dangerous form of bias that may obscure other underlying root causes.

The application of non-compliance analysis is further complicated by the difficulty of determining whether or not particular regulations and policy documents are applicable to particular projects. This might seem to be a trivial task in many industries. However, NASA preferred practice procedures were drafted by individual centres during the period preceding the loss of the Polar Lander and the Climate Orbiter. For example, Practice No. 1437 on end-to-end compatibility and mission

simulation testing explicitly states that “all flight programs managed by the Goddard Space Flight Center (GSFC) are required to use this practice” [568]. This situation is not uncommon. Different regional or function divisions often draft supplementary regulations to support their particular activities. Problems arise when investigators must determine whether local regulations affected the course of an incident and whether they interacted with the requirements that are imposed at other levels within an organisation or from regulatory organisations.

The individual nature of many NASA projects can prevent investigators from establishing the norms that govern development and operation practices. Each project is so different that it can be difficult to identify which of those differences actually contributed to an incident. This makes it difficult for investigators to use techniques, including change analysis, that focus on the differences between ‘normal’ and observed behaviour. Non-compliance analysis suffers from similar problems. Differences between projects force managers to adapt existing working practices. For instance, radical changes in the relationships between JPL, NASA Headquarters and the subcontractor organisations forced program managers to adapt existing reporting procedures during the Mars Surveyor’98 program. They also complicate any attempts to enumerate those policies and regulations that govern each stage of the missions within each of the participant organisations. NASA recognise the need for flexibility in the face of changing mission demands. For instance, NASA Standard 8729.1 is one of several guidelines that specifically allows departures from the recommended practice. Such flexibility creates difficulties for investigators who must determine whether or not it was reasonable for projects to decide not to comply with recommended practice:

“Section 1.3 Approval of Departures from this Standard. This standard provides guidance and is not intended for use as a mandatory requirement; therefore, there is no approval required for departing from this standard. However, the fundamental principles related to designing-in Reliability and Maintainability (R&M), as described in this standard, are considered an integral part of the systems engineering process and the ultimate R&M performance of the program/project is subject to assessment during each of the program/project subprocesses (Formulation, Approval, Implementation, and Evaluation).

A third factor that complicates non-compliance analysis is that there may be genuine uncertainty within an organisation about whether or not an individual should have complied with particular regulations. This is apparent in JPL’s response to the Faster, Better, Cheaper strategy. This initiative led individual managers to reassess whether or not particular policies, for instance concerning the use of model-based validation rather than destructive testing, were still appropriate to the new context of operation:

“(NASA, JPL and LMA) have not documented the policies and procedures that make up their FBC approach; therefore, the process is not repeatable. Rather, project managers have their own and sometimes different interpretations. This can result in missing important steps and keeping lessons learned from others who could benefit from them. [570]”

It is relatively easy in retrospect to argue that an incident occurred, therefore, a regulation was violated. It is less easy to determine whether any individuals within the organisation would have concurred with that analysis *before* the incident took place. This hindsight bias is a particular danger where non-compliance analysis is (ab)used as a mechanism for blame attribution.

It can also be difficult to apply compliance analysis to the results from previous stages in a causal factor analysis. For instance, the following list enumerate the causal factors that were identified for the Deep Space 2 and Polar Lander mishaps. These causal factors were derived by applying counterfactual reasoning to each of the events that was represented within previous ECF charts of this incident:

1. Both DS2 probes suffer electrical failure at impact
2. Forces at impact compromise aft body battery assembly

3. Forces at impact compromise RF components
4. Premature Shut-Down of engines (40 meters above surface)
5. Software marks individual legs as failed if they show spurious signals but does not reset touch-down indicator at 40 meters (entry +5:16)
6. Transient signals possible from Hall Effect magnets when legs first deploy at 1,500 meters (Entry +4:13)
7. Launch approved
8. Preliminary Design Review passed

It is difficult to directly apply non-compliance analysis to any of these causal factors. For example, the electrical failure of the Deep Space 2 probes on impact cannot itself be blamed upon a lack of knowledge about applicable regulations or on an inability to meet those regulations or on a deliberate failure to follow those regulations. This is because the causal factor related to a direct failure rather than to any particular form of non-compliance by an identifiable individual or group. A further stage of analysis is required before investigators can exploit this categorisation as a means of identifying potential root causes. For instance, the failure of Radio Frequency components on impact with the planet surface is a probable failure mode because development impact tests were limited to brassboard and breadboard components and subassemblies [580]. Visual inspections were conducted after these test to ensure that the component mountings and the electrical connections remained intact. Unfortunately, many of the components were not electrically functional. As a result, it was only possible to conduct limited inspections of the powered circuits before and after the impact tests. In other words, the impact tests established the structural integrity of the design but did not establish the functional validity. It can, therefore, be argued that the RF testing during the development of the Polar Lander indicates non-compliance with NASA requirements. In particular, Preferred Reliability Practice PT-TE-1435 governed the verification of RF hardware within JPL from February 1996. Impact tests are implied by a requirement to evaluate RF subsystem performance under ‘other environmental conditions’:

“Analyses are performed early in the design of radio frequency (RF) hardware to determine hardware imposed limitations which affect radio performance. These limitations include distortion, bandwidth constraints, transfer function non-linearity, non-zero rise and fall transition time, and signal-to-noise ratio (SNR) degradation. The effects of these hardware performance impediments are measured and recorded. Performance evaluation is a reliability concern because RF hardware performance is sensitive to thermal and other environmental conditions, and reliability testing is constrained by RF temperature limitations.” [579]

The failure to follow PT-TE-1435 is classified as an inability to comply. It is, therefore, associated with root causes that centre on resource allocation issues. This judgement is supported by the finding that there were several design changes late in the development program that prevented impact testing without jeopardising the launch of the Polar Lander. If the battery cells and RF subsystem assemblies had been available earlier in the development cycle then it might have been possible to comply with PT-TE-1435. This line of analysis is summarised by the non-compliance diagram illustrated in Table 10.23.

If we continue this non-compliance analysis, the situation is shown to be considerably more complex than that suggested in Table 10.23. In particular, the Preferred Practice proposed in PT-TE-1435 centres on the use of modelling as a means of validating the initial design of RF components. This is particularly important because mathematical analysis can be used to identify potential design weaknesses before projects accept the costs associated with procuring particular subsystems. PT-TE-1435 argues that these models help in situations where it is “difficult to pinpoint the exact cause of unexpected test results once the subsystem has been integrated”. [579] From this it follows that the development team could have complied with PT-TE-1435 even though design changes meant

Causal Factor	Procedure or Regulation	Compliance Failure?
Forces at impact compromise RF components	Preferred Reliability Practice PT-TE-1435 Early validation of RF reliability under thermal and other environmental conditions.	Can't comply RF assembly unavailable for impact testing as design changes delay development.

Table 10.23: Non-Compliance Analysis of RF Failure Mode on Deep Space 2 Probe.

that the flight unit was not available for impact tests. Mathematical models could have been used to provide the validation recommended in this regulation. Unfortunately, the impact analysis of high gravitational forces does not yield reliable results. Finite element analysis was used to validate the antenna structure. This did not provide reliable results because the impact loads were not well understood. Several antenna masts were slightly bent during impact testing, but no analytic models could be made to match the empirical damage. Empirical impact testing provides the only reliable verification method.

As before, further analysis of this apparent non-compliance can yield further insights into the complexities that characterised the development and testing of the Deep Space 2 probes. NASA requirements, such as PT-TE-1435, were well understood by JPL employees and the contractor organisations. The design changes to the RF system meant that any impact tests would not be completed before the scheduled launch of the Polar Lander. They, therefore, attempted to gain explicit approval for the decision to proceed to launch without an RF subsystem impact test:

“The DS2 project thought there was no alternative to accepting the absence of a flight-like RF Subsystem impact test, short of missing the MPL launch opportunity. The rationale for proceeding to launch was presented and accepted at two peer reviews and presented at three project-level reviews: Risk Assessment, Mission Readiness, and Delta Mission Readiness. The project had proceed to launch concurrence from JPL and NASA upper management.” [580]

Such actions can be interpreted as an understandable reluctance to comply with the requirements and recommended practices that governed RF validation. Mission schedule was interpreted within the Faster, Better, Cheaper strategy as being more critical than additional reliability tests for components that had already been validated at a structural and component level. Table 10.24, therefore, builds upon the previous analysis to document these additional reasons for non-compliance.

The initial resource allocation problems, connected with late design changes to RF components, were compounded by the pressures to launch on schedule. Higher-levels of management were prepared to concur with this decision, arguably, because of the perceived need to implement the the Faster, Better, Cheaper strategy. This illustrates the way in which non-compliance analysis helps to identify the deeper root causes of an incident. The specific causal factor revealed by the causal factor analysis is unlikely to threaten future missions simply because it has been identified as a potential cause of the Deep Space 2 mishap. The validation of RF assemblies will include system-level impact tests. In contrast, the root cause of the non-compliance remains a concern for subsequent missions. Mission deadlines and tight launch schedules will continue to encourage engineers and managers to sanction non-compliance with accepted working practices. The mishap report into the management structures that contributed to the loss of the Climate Orbiter observed that:

“NASA currently has a significant infrastructure of processes and requirements in place to enable robust program and project management, beginning with the capstone document: NASA Procedures and Guidelines 7120.5. To illustrate the sheer volume of these processes and requirements, a partial listing is provided in Appendix D. Many of these clearly have a direct bearing on mission success. This Boards review of recent project failures and successes raises questions concerning the implementation and adequacy of existing processes and requirements. If NASA programs and projects had

Causal Factor	Procedure or Regulation	Compliance Failure?
Forces at impact compromise RF components	Preferred Reliability Practice PT-TE-1435 Early validation of RF reliability under thermal and other environmental conditions.	Can't comply 1. RF assembly unavailable for impact testing as design changes delay development. 2. Mathematical modelling of high g impacts yields unreliable results. Won't comply 1. JPL and NASA upper management approve launch without RF impact validation in order for DS2 to meet launch schedule. 2. RF subsystem components had been structurally tested and were similar to other components used in previous missions.

Table 10.24: Non-Compliance Analysis of RF Failure Mode on Deep Space 2 Probe (2).

implemented these processes in a disciplined manner, we might not have had the number of mission failures that have occurred in the recent past.” [570]

The Appendix of the report lists over fifty NASA standards that were identified as relevant to this incident. These ranged from standards relating to electrical discharge control through safety-critical software development to standards for oxygen systems. This not only reflects the complexity of any non-compliance analysis, mentioned above, but it also illustrates the demands that are place on managers and operators who must ensure compliance to these regulations while also satisfying high-level mission objectives such as those implied by the Faster, Better, Cheaper strategy.

10.4 Summary

This chapter has shown how a range of diverse analytical techniques can be used to identify the causal factors that contribute to a particular incident. These causal factors can then be used to determine the underlying root causes that might continue to threaten the safety of future systems. The techniques that we have exploited are based on those advocated by the US Department of Energy. Their approach was specifically developed to support the analysis of workplace injuries. It has not been widely applied to reason about the causes of complex, technological failures. This is surprising given that NASA’s Procedures and Guidelines document NPG:8621.1 on mishap reporting recommends this same approach to root cause analysis. We, therefore, demonstrated that these techniques could be used to support an investigation into the loss of the Mars Climate Orbiter and the Mars Polar Lander missions. These case studies are not ‘safety-critical’ in the sense that they did not threaten human life after they had left the Earth’s orbit. They do, however, reflect a more general class of mission-critical incidents that are considered by many reporting systems. These case studies were also chosen because they provide an extreme example of the technological complexity and coupling that characterises many safety-critical failures. The Climate Orbiter and

Polar Lander missions also provide a strong contrast with the level of technology involved in the Allentown explosion in Chapter 8.3.

This chapter began with the construction of ECF charts. These graphs help to identify the events and conditions that lead to an incident. They are similar to modelling techniques, especially graphical time-lines and Fault Trees, that have been introduced in previous chapters. They do, however, suffer from a number of potential limitations. In particular, ECF charts can bias investigators towards the representation of observable events rather than the wider contextual factors that made those events more likely. The US Department of Energy guidelines and the NASA procedures advocate the use of supplementary analytical techniques to uncover these factors. For instance, change analysis can be used to identify the impact that different management priorities, new working practices and technological innovation have upon the course of an incident. These changes often lead to the unanticipated interactions that have been identified as important causes of ‘systemic’ failures [486]. Similarly, barrier analysis helps to move the focus away from events that actively contribute to an incident. This technique encourages investigators to consider the ways in which a wide variety of potential barriers must fail in order for an incident to occur. Both of these analytical techniques can be used to look beyond the initial events that are represented in an ECF chart. They encourage investigators to revise those diagrams and, in particular, to incorporate a wider range of causal factors.

The causal factors are distinguished from a wider range of contextual factors using causal factor analysis. This technique involves the use of counterfactual reasoning. For each event in the revised ECF chart, investigators must ask ‘would the incident have occurred without this event?’. If the answer is yes then the event is not considered to be a causal factor. If the answer is no then investigators must record further information about the event. This information centres on a number of prompts including: what led to the event? What went wrong? How did the barriers fail? Who was involved in the event? Is the event linked to a more general deficiency? The results of this more detailed analysis can be recorded in an ECF table. These, in turn, are used to drive any subsequent root cause analysis.

Causal factors are identified using counterfactual reasoning. An incident would not have occurred, if the event or condition had not occurred. In contrast, root causes are events or conditions that threaten the safety of future systems. They often result from the amalgamation of several causal factors. For example, the failure of several barriers may indicate a more general failure to ensure adequate protection. Any attempt to fix particular barriers will still leave a concern that other barriers may still be susceptible to other forms of failure until this root cause is more directly addressed. Several techniques have been proposed to help investigators move from specific causal factors to these more general root causes. Again our use of tier and non-compliance analysis has been guided by the US Department of Energy’s recommendation. Tier analysis depends upon the development of tables that associate causal factors with different levels in an organisational structure. The entries in these tables are then inspected in order to identify more general patterns that might indicate a root cause that is common to several causal factors. In contrast, non-compliance analysis involves the examination of any rules or procedures that might have been violated either directly by an event or by the wider conditions that made an event more likely.

It is important to emphasise that the techniques which we have described do not provide a panacea for the problems of root cause analysis. It can be difficult to apply some of these approaches to the specific circumstances that characterise particular technological failures. The documentation techniques that are associated with key stages in the analysis, especially the revised ECF charts, are cumbersome and intractable. All of the techniques that we have described rely upon the subjective skill and experience of individual investigators. The insights that they provide must, therefore, be validated by other members of an investigation team or a safety management group. A number of researchers are currently working to produce automated systems that remove some of the subjectivity involved in root cause analysis. Unfortunately, sophisticated reasoning tools often impose unacceptable constraints upon the way in which an incident is modelled. The syntax and semantics of any input must be narrowly defined so that the system can recognise and manipulate model components during any subsequent root cause analysis. There are a number of potential solutions to this problem, including structural induction over graphical structures similar to ECF chart. In anticipation of

the results of this research, it is difficult to underestimate the importance of the tables and diagrams that are presented in this chapter. They provide other analysts and investigators with means of tracing the reasons why particular events and conditions are identified as causal factors. They also help to document the process by which root causes are determined. Without such documents, it would be extremely difficult to validate the subjective analysis of incident investigators.

The penultimate remarks in the Chapter belong to Daniel Goldin; the NASA Administrator who first formulated the Faster, Better, Cheaper strategy. He spoke to the engineers and managers at the Jet Propulsion Laboratory about the loss of the Climate Orbiter and the Polar Lander.

“I told them that in my effort to empower people, I pushed too hard... and in so doing, stretched the system too thin. It wasn’t intentional. It wasn’t malicious. I believed in the vision... but it may have made failure inevitable. I wanted to demonstrate to the world that we could do things much better than anyone else. And you delivered – you delivered with Mars Pathfinder... With Mars Global Surveyor... With Deep Space 1. We pushed the boundaries like never before... and had not yet reached what we thought was the limit. Not until Mars 98. I salute that team’s courage and conviction. And make no mistake: they need not apologise to anyone. They did not fail alone. As the head of NASA, I accept the responsibility. If anything, the system failed them.” [575]

There is a danger that the recent emphasis on systemic failures will discourage investigators from pursuing the coherent analysis of specific root causes. Many incidents are characterised by emergent behaviours that stem from complex interactions between management practices, operational procedures and particular technologies. These interactions are not, however, random. They are shaped and directed by the regulatory environment and by higher-levels of management. Goldin’s words are important because they acknowledge personal and corporate responsibility for the systemic factors that led to failure.




Figure 10.12: Technological Barriers Fail to Protect the Climate Orbiter




Figure 10.13: Technological Barriers Fail to Protect the Climate Orbiter (2)




Figure 10.14: Integrating Change Analysis into an ECF Chart




Figure 10.15: Representing Staffing Limitations within an ECF Chart




Figure 10.16: Representing Risk Management Issues within an ECF Chart




Figure 10.17: Representing Technological Issues within an ECF chart (1)




Figure 10.18: Representing Technological Issues within an ECF chart (2)



Figure 10.19: Using Change Analysis to Collate Contextual Conditions




Figure 10.20: Integrating Development Issues into an ECF chart (1)




Figure 10.21: Integrating Development Issues into an ECF chart (2)




Figure 10.22: Integrating Review Issues into an ECF chart




Figure 10.23: An ECF chart of the Deep Space 2 Mission Failure



Figure 10.24: An ECF chart of the Polar Lander Mission Failure




Figure 10.25: An ECF chart of the Climate Orbiter Mission Failure




Figure 10.26: NASA Headquarters' Office of Space Science [570]




Figure 10.27: JPL Space and Earth Sciences Programmes Directorate [570]