Chapter 8

Computer-Based Simulation

The previous chapter identified the main activities that must be conducted during the secondary investigation of any incident report. These, typically, focus on gathering evidence to both inform and validate initial hypotheses about the causes of an adverse occurrence. This chapter, in contrast, looks at one aspect of this validation process. It seems clear that any causal hypothesis must be consistent with what we know about the course of any incident. Support for such assertions is often provided by simulation and reconstruction techniques. The Rand report into the National Transportation Safety Board's (NTSB) investigation techniques emphasised this in their overview:

> "When a complex system fails, the number of potential scenarios rises proportionately. NTSB investigators must carefully unravel the performance of many highly integrated systems, a very time-consuming task requiring a diverse set of skills. Often, this requires extensive and costly salvage and reconstruction of the aircraft. Complexity affects more than just staff workload. The growing complexity of aircraft crashes also has a profound effect on how investigations must be structured to reveal hidden failure modes." [482]

This quotation reveals the dual nature of reconstruction in many modern incident investigations. Firstly, reconstruction involves the rebuilding of components and sub-components to identify causal information from the physical damage that often occurs during major incidents. Secondly, there is the more abstract notion of event simulation in which investigators piece together the more complex causes of an incident drawing upon the physical evidence and also from the other forms of evidence that are gathered during a secondary investigation. This might suggest a firm distinction between physical reconstructions and virtual simulations. The terms 'simulation' and 'reconstruction' are, however, often used inter-changeably. As we shall see, this ambiguity can partly be justified by the way in which limited physical reconstructions are being used to provide the data that drives more general computer-based simulations.

8.1 Why Bother with Reconstruction?

The term 'reconstruction' has traditionally been used to describe the way in which physical evidence is re-assembled to provide clues about the sequence of events leading to failure. For example, the US Army's accident and incident investigation guidelines constraint the following recommendations for the analysis of rotor or propeller failures. As can be seen there is a requirement to reconstruct the entire assembly if at all possible:

1. Collect and inventory; reconstruct the whole assembly if possible.
2. Examine damage / scarring to determine if systems were turning at impact and if power was applied at impact.
3. Examine all linkage from cockpit controls to systems for continuity/disconnect, all bearing assemblies and / or blade grips for failure prior to impact.

4. Check for serial numbers of blades / propellers against historical records. [807]

A number of published guidelines provide detailed information about the ways in which such physical reconstruction should be conducted [751]. Much of this information varies from domain to domain. For instance, the construction and operational stresses of aircraft components are quite different from those relating to automotive components. Klepacki, Morin and Schaeffer's guidelines for evaluating post-incident flight control trim system configurations are highly domain specific [448]. There are, however, some similarities between the techniques that are used to support incident reconstruction in several different industries. For instance, the techniques for establishing the velocity and angle of impact damage show strong similarities across several different domains [870].

The opening paragraphs identified two forms of reconstruction. The first focussed on the physical rebuilding of damaged components to gain further information about the failures that contributed to an incident or forces that arose in its aftermath. The second aspect of reconstruction deals with the way in which information is used to describe the course of events over time. This centres on the process of assembling fragmentary evidence to produce a coherent account of an adverse occurrence. For example, the US Air Force requires that investigators reconstruct the sequence of events that leads to an incident [795]. They must map route segments. They should provide a vertical view of maneuvers. The account may include artists conceptions or models to explain the course of events. The intention is to "explain what the plan was, what should have happened if things had gone right, who was in charge, what were the rules of engagement and were they followed, where things went wrong, what should the aircrew have done, and what were the aircraft parameters at ejection or aircraft impact" [795]. As can be seen, generic requirements are specified together with more domain specific guidelines that relate narrowly to aviation accidents.

The wealth of guidance on the physical reconstruction of safety-critical systems is not matched by similar sources of advice on the reconstruction of events leading to incidents and accidents. As a result, the remainder of this chapter focuses on techniques that can be used to build coherent models that explain how different events contribute to an adverse occurrence. It is important not to underestimate the importance of these reconstructions. They are intended to produce a coherent account of the course of an incident from many disparate pieces of evidence. In other words, they are intended to explain *what* happened while causal techniques present *why* it happened.

There are many different ways in which to build these event reconstructions. In some domains, it is also possible to stage physical reconstructions. These re-enactments are often used by the Police to trigger witness recollections and elicit further information about an incident. However, there are obvious limitations with this approach. For example, such reconstructions can expose individuals to further danger. There are ethical considerations involved in re-enacting a failure in a working foundry or chemical plant. There is also the danger, described in Chapter 4.3 that such 'realistic' reconstructions may trigger further psychological problems for those involved in an incident. It may even trigger false memory syndrome in some cases. Fortunately, a range of alternative techniques can also support the reconstruction of safety-critical incidents. For example, computer-based simulations enable investigators to step-through the events leading to an incident using three dimensional animations. The second half of this chapter identifies a number of these interactive tools. There are, however, a number of limitations with computer-based simulations. For example, highly interactive models provide a good impression of the events leading to component failure. Unfortunately, they cannot easily be used to recreate the events leading to managerial or regulatory failure. I have yet to see virtual reality simulations recreate a board meeting or a management conference in which safety investments were turned down!

Fortunately, a range of graphical and textual notations can be used to avoid such limitations. They can be used to sketch the events leading to an incident at a far greater level of abstraction and can, therefore, also capture events leading to managerial and regulatory failure. Many of these notations provide inference techniques that can be used by investigators to clearly distinguish what can, and what cannot, be concluded from the evidence that is assembled. These reasoning techniques can also be used to identify inconsistencies and omissions in incident reconstructions. The following chapter focuses on these formal and semi formal notations for event reconstruction. The second half of this chapter reviews computer-based modelling techniques. It is important to emphasise, however, that these modelling notations and the computer-based simulations mentioned above are *not* primarily intended as tools to support the causal analysis of adverse occurrences. They are intended to reconstruct the course of events that contribute to an incident. In contrast, Chapter 9.3 presents techniques that help to distinguish root causes and contributory factors from the contextual information that can be represented in a formal model or an interactive simulation.

Some formal and semi-formal modelling techniques are supported by computer-based tools. It is possible to derive interactive computer-based simulations from abstract notations. As a result, the previous distinctions between computer-based simulation and abstract modelling can become blurred. However, these distinctions are retained because they are useful in distinguishing between different sets of concerns that affect what can be complementary approaches. For instance, formal mathematically-based notations can be difficult to interpret by non-mathematicians. In contrast, interactive simulations can often lead to unwarranted interpretations if analysts are seduced by particular animation techniques. For now it is sufficient to emphasise that natural language provides the most accessible and widespread medium for building reconstructions. The following quotation illustrates how most organisations summarise the events leading to an incident These accounts are used to provide investigators and managers with a common ppoint of reference during any investigation. They are gradually refined as additional evidence is obtained until they are eventually integrated into a final report:

> "About 10 p.m., unknown to the controller, the pipeline ruptured at a location near Gramercy, Louisiana. At 10:01:53 p.m., the supervisory control and data acquisition (SCADA2) system reported high-pump-case pressure at Garyville. The SCADA system activated an audible alarm and also displayed a message on a display screen. Almost immediately, the SCADA system sounded and displayed alarms reporting that certain pumping units at the Garyville station had automatically shut down because of low suction pressure (low liquid pressure on the inlet side of the pump). At 10:02:30 p.m., the SCADA system reported a line balance alarm." [595]

This quotation shows the way in which natural language can represent the events that occurred immediately before the rupture. The model of this incident is constructed around a number of key incidents, such as the first SCADA2 system report, for which the timing is known. It is important to note that these proximal events cannot be viewed as catalytic because no argument is provided to demonstrate that they directly *caused* the incident. Of course, it is also possible to use natural language to describe the more distal, latent causes of this failure. This again illustrates how the reconstruction of an incident is guided by implicit causal hypotheses. In this case, evidence about previous excavations near the Marathon pipeline suggests that damage might have been caused during these earlier operations:

> "The investigation determined that in 1995, LaRoche Industries, Inc., arranged for excavation of and repairs to various portions of its 8-inch pipeline, which was located about 30 feet from the Marathon pipeline. These excavations took place in September and October 1995 in the vicinity of the Marathon pipeline rupture... According to officials from LaRoche's contractor, the equipment operators were told by LaRoche superintendents that no pipelines were located in the area of the Marathon pipeline." [595]

The previous quotations illustrate how prose descriptions can be used to draw upon various sources of evidence in order to reconstruct the events that led to an adverse occurrence. These natural language accounts must consider the many different factors that contribute to the increasingly complex incidents that have been described in previous chapters. The following paragraphs briefly summarise the types of information that must be captured. These can be loosely categorised as belonging to three distinct stages in an 'incident sequence'.

Initial Conditions

The initial conditions describe the normal operating state of the system and its environment before an incident occurs. The following quotation illustrates how NTSB investigators describe the initial state of the system as part of a passage that sets the scene for the failures that follow:

> "On May 23, 1996, a pipeline controller was on duty in Marathon Pipe Line Company's pipeline operations center in Findlay, Ohio, operating and monitoring a 68-mile-long segment of Marathon pipeline located in Louisiana. This pipeline is used to transport hazardous liquids between a refinery at Garyville, Louisiana, and a station at Zachary, Louisiana. Pumps at the Garyville refinery pressurise the pipeline and generate the power to transport the liquids to the Zachary station. About 9:53 p.m. central daylight time on May 23, the pipeline controller had just completed operations to transport a batch of unleaded gasoline through the pipeline. He then remotely executed commands to introduce into the pipeline (behind the gasoline) a batch of 125,000 barrels of low-sulfur diesel fuel." [595]

This excerpt establishes the general topology of the pipe network. It also introduces the controllers' tasks immediately before the incident took place. As can be seen, the initial conditions in a reconstruction describe the situation as it existed before any adverse incidents occurred. This introductory section closes at the moment when the remote command was executed. At this point the pipeline ruptures and the reconstruction continues with the first of the quotations cited in the previous section.

It is important to emphasise that the initial conditions that are described by a reconstruction need not be normative. They may not satisfy relevant safety regulations or recommended operating practices. In particular, the initial description of the system might indicate that there were frequent safety violations during normal operation. For example, the NTSB overview of another pipe rupture describes how the company "had procedures in place at the time of the accident that were applicable to general construction activities in proximity to its pipelines, but it did not have procedures specific to directional drilling operations' [599]. These operating practices did not have adverse consequences until the company attempted to install a distribution main parallel to a gas transmission pipe. The proximity of the installation damaged the transmission pipe and this led to a rupture that the NTSB estimates cost in excess of $2 million.

Catalytic Failures

Reconstructions must also describe the events that helped to move the system from its initial condition towards an eventual incident. These events include the catalytic failures that are often central to any subsequent investigation. In many instances these are clear cut. For example, some pipeline ruptures can be directly related to specific (catastrophic) events:

> "About 4:50 a.m. on October 23, 1996, in Tiger Pass, Louisiana, the crew of a dredge dropped a stern spud into the bottom of the channel in preparation for dredging operations. The spud struck and ruptured a 12-inch-diameter submerged natural gas steel pipeline owned by Tennessee Gas Pipeline Company. The pressurised natural gas released from the pipeline enveloped the stern of the dredge and an accompanying tug, then ignited, destroying the dredge and the tug." [594]

It is important again to reinforce the point that such events do not provide any direct explanation of the root causes of an incident. Such underlying causes are often embedded within the initial conditions that were mentioned in the previous paragraph. These conditions combine to create a situation in which catalytic events have the potential to trigger an incident or accident. For example, the dredging operation had to rely upon the gas company's practices and procedures for locating, marking, and maintaining markers for gas pipelines through navigable waterways. The potential for a catalytic event was also created by the lack of Federal requirements for placing and maintaining permanent markers where gas pipelines cross navigable waterways.

It should also be emphasised that there are other incidents in which it is far harder to identify the catalytic events that actually trigger a failure. For instance, in the incident in which directional drilling was cited as the root cause of the pipeline fracture, it is unclear as to which aspect of the operation actually *caused* the failure. The rupture occurred as the distribution pipeline was being returned to full service and not as an immediate consequence of a catastrophic operation as was the case in the dredging incident. Some evidence pointed to the fact that a reaming tool left gouge marks in the vicinity of the rupture but it is difficult to be certain of the precise operation that resulted in the eventual failure. In such circumstances, there must be some ambiguity or uncertainty in an eventual reconstruction of the catalytic events that contributed to an incident.

Liveness Conditions

There are, typically, several moments when operators or automated systems might have intervened in order to prevent an incident. 'Liveness' conditions, therefore, not only describe the catalytic events introduced in the previous paragraph. They also reconstruct the manner in which these safeguards failed. In other words, liveness conditions also describe the events that enabled an incident to progress towards its final consequences. For instance, the NTSB account of the gasoline release describes how operators failed to detect and respond to the initial alarms from the SCADA system. Had they responded to these warnings sooner then the full consequences of the incident might have been considerably reduced:

> "The pipeline controller continued to receive alarms. Initially, he acknowledged each one individually, but believing that each subsequent alarm was related to operations at the refinery, he elected to simultaneously acknowledge all the alarms and the alarm text messages without attending to the nature of each alarm... The controller said he called Garyville and discussed the situation with the station operator there. The station operator confirmed the automatic pump shutdowns. The station operator determined that the Garyville refinery was, indeed, loading product to a barge. Even though refinery personnel reported that the volume of product being delivered was insufficient to have caused the SCADA system to alarm, the pipeline controller and the station operator concluded that the loading of the barge had precipitated the alarms and the pump shutdowns." [595]

To summarise, reconstructions must represent the initial conditions or context for an incident. They must also describe the way in which catalytic failures initiate the events leading to failure and how liveness conditions create the necessary conditions for an incident to develop. The following section describes the final component of any reconstruction; the events that take place in the aftermath of an adverse occurrence.

Consequences

The reconstruction of an incident cannot simply stop at the point with catalytic events. It is a truism that more lives are lost through failures in the 'golden hour' after an accident has occurred than are killed by the catalytic events themselves. Chapter 6.4 has emphasised the importance of incident reporting as a means of assessing an organisations ability to respond to or mitigate the adverse consequences of any failure. In consequences, reconstructions must go beyond the catalytic events that are often the focus of attention in the aftermath of any incident. This point can be reinforced by the consequences of two further pipeline failures. On October 30th 1998, excavation work damaged a 24-inch diameter gas main in Chicago. This released natural gas that ignited about forty minutes after the initial rupture. The immediate consequences of the failure included substantial damage to a high-rise block of appartments. However, the prompt response to fire and police personnel completely evacuated the building so that no-one was injured. In contrast, a similar incident two months later left four dead, one person seriously injured and ten people, including two firefighters and a police officer, with minor injuries:

> "An engine company with a lieutenant and three firefighters arrived within minutes of fire department notification. Firefighters attempted to take gas concentration readings with a gas monitor, but the monitor had not been calibrated in fresh air and gave invalid or unreliable readings. Firefighters continued to attempt readings with the improperly calibrated instrument, all the while working in an environment in which they described the gas smell as pretty bad. At no point did firefighters check buildings near the leak site to determine if natural gas was accumulating or to help assess the need for

> a possible evacuation, even though the gas line was continuing to release gas that could migrate through the ground and into nearby buildings, where it could present a danger of explosion. Two of the firefighters near the leak site returned to their truck as soon as two gas company employees arrived. It should have been obvious to the firefighters that a threat continued to exist and that the situation could worsen. The Safety Board therefore concludes that firefighters of the St. Cloud Fire Department responded quickly to the scene of the leak; however, once on the scene, the firefighters actions did not fully address the risk to people and property posed by the leak or reduce the consequences of a possible fire or explosion." [604]

Many investigation authorities have placed increasing emphasis on response time targets for emergency services. This incident again illustrates that a prompt response must be backed up by effective actions if we are to mitigate the effects of such incidents. It is, therefore, necessary for reconstructions to explicitly represent the actions that are taken in the aftermath of a catalytic failure.

Previous paragraphs have explained how it is important to reconstruct the initial conditions that create the context for any incident. It has also been argued that investigators must produce a coherent account of the catalytic failures that trigger those events that lead to an adverse occurrence. Liveness conditions must also be reconstructed. These represent the way in which defences must be breached and warnings ignored in order for a catalytic event to escalate into a major failure. Finally, it has been argued that reconstructions must also consider the consequences of any incident. These are partly shaped by the nature of the failure but also by the interventions that help to mitigate those consequences. However, there are relatively few benefits to be obtained from simply developing accounts that describe how all of these events occurred during the course of an incident. The following paragraphs, therefore, identify ways in which reconstructions can be used to inform the investigation of an adverse occurrence and, ultimately, to reduce the likelihood of any recurrence.

8.1.1 Coordination

Chapter 6.4 has described how incident investigations draw upon the the work of many different experts. Forensic scientists, metallurgists, meteorologists, software and systems engineers as well as human factors experts all contribute to these enquiries. It can be difficult to coordinate the activities of these different group. There is a risk that necessary tasks may be omitted or needlessly duplicated. It is, therefore, important that investigators have some means of monitoring and coordinating the finite resources that they can deploy to support their enquiries into an incident. Reconstructions provide a useful tool to support these managerial tasks. They provide a model of the events leading to an incident. Individuals with different domain expertise contribute to different aspects of these reconstruction. For example, metallurgists can describe the conditions that might have contributed to catalytic metal fatigue. Human factors experts can identify salient events in a crews' response to an incident. These different contributions must be pieced together to form a coherent view of a complex incident.

There are a number of ways in which experts contribute to the overall process of incident investigation through their participation in any reconstruction. These can be summarised as follows:

- *broadening the causal field.* One of the key roles for any expert is to help broaden the causal field of any analysis. Chapter 6.4 explained how this field represents a subjective frame of reference that individuals or organisations use when trying to explain what has happened in a particular situation. If an event does not have an impact upon the causal field then it may not be identified as playing a significant role in the course of an incident. Prior expertise plays a significant role in knowing where to find the evidence that indicates certain events have taken place. Without this expertise, evidence might not be found and a reconstruction might not include necessary information about an incident. .

- *determining the salience of events.* In contrast to the experts' role in broadening a causal field, they may also identify certain events as not playing a significant role in a particular incident. These events might then be omitted from any subsequent reconstruction. There is, of course, a potential danger in this if those events later emerge as having a more important role in course of events. It is, therefore, important to document the reasons for such omissions. It is also important to stress that reconstructions support but do not replace causal analysis. For example, it is often impossible for any single expert to diagnose the root cause of an incident without referring to the work of their colleagues. In consequence, investigators use reconstructions to provide an overview of the evidence that is collected about the diverse events that lead to an incident. This overview of events must then be interpreted and analysed to distinguish between contextual factors, contributory factors and root causes. Techniques that support this causal analysis will be described in the next chapter. In contrast, this chapter focuses on techniques that can be used to reconstruct the 'flow' of events leading to an incident and the consequences that stem from such failures.

- *determining knock-on effects of events.* By participating in any reconstruction, experts are also forced to consider the ways in which events in other areas of a system can affect their area of expertise. For instance, a human factors expert must consider the impact of prevailing weather conditions if a meteorological experts have indicated that this may be a factor. Conversely, building a reconstruction can also help investigators to identify the knock-on consequences that particular events will have throughout a system. This is a by-product to the tasks involved in developing a narrative account that links together the evidence that is available in the aftermath of an incident.

- *eliminating particular events.* The development of a reconstruction can force analysts to determine whether or not particular events actually did contribute to an incident. For example, the Minnesota pipeline investigation initially questioned whether the location of the line had been incorrectly marked out. Evidence had to be provided to determine whether this was a potential problem before any detailed reconstruction could be built. If it had been incorrectly marked out then additional resources would have been deployed to examining the events that contributed to this failure. However, the NTSB investigation determined that "the marked location of the ruptured gas line was accurate and therefore not a factor in this incident" [604]. As a result, the prose description of the incident does not focus in great detail on the initial surveying of the line.

- *forcing the resolution of inconsistency.* As investigators contribute to the development of a reconstruction, it is likely that a number of inconsistencies and omissions will be identified in the overall timing of events. As we shall see, other anomalies can arise. For example, eye-witness testimonies often place the same individual at two different locations at the same moment in time. Such inconsistencies can be resolved by finding evidence to discount one statement. Alternatively, two or more alternative reconstructions can be developed to explore several different incident scenarios. However, contradictory witness statements are not the only source of inconsistency in incident reconstructions. Other problems relate to the group processes that affect incident investigation teams. As mentioned, these enquiries often involve heterogeneous teams of domain experts. The members of these groups often have different backgrounds and training. Partly as a result of this, the conclusions of one analyst about the probable ordering of events need not accord with those of another. For example, the Fire Service and Ambulance accounts of the Clapham rail crash differ in several important respects [502]. If these problems are not resolved during the reconstruction phase then there is a danger that any causal analysis will be jeopardised because of contradictions in the evidence that it relies on. There is also the danger that the eventual incident report will contain inconsistent information about the sequencing of events leading to and stemming from catalytic failures.

A number of important consequences stem from this use of reconstructions as a means of coordinating the various activities that contribute to an incident investigation. In particular, natural language descriptions, interactive computer simulations or other diagrammatic techniques must be capable of capturing the key events identified by different domain experts. It is also important that those domain experts can read and understand the resulting reconstructions if they are to validate the models that are produced. This is a non-trivial requirement. The complexity of many incidents makes it difficult to trace the ways in which system 'failure' and operator 'error' interact over time. For example, many incident reports now run to several hundred pages of prose narrative.

8.1.2 Generalisation

Incident investigations are intended to determine what caused a failure and to identify means of preventing any recurrence. As we have seen, reconstructions play an important role in validating the evidence that, in turn, supports subsequent causal analyses. They also play an important role in identifying ways in which an incident can recur. The process of identifying those events that contributed to a particular incident helps to inform subsequent investigations to determine whether those events might recur in isolation or in combination with other failures. In other words, in order to identify the general causes of future incidents it is important to understand the causes of the particular incident under investigation [678]. Those causes can only be accurately established by ensuring the validity of any reconstruction.

An important application of reconstruction techniques is in the development of training scenarios. These, typically, start with the events that lead to previous failures. For example, the following citation comes from an NTSB incident report that explicitly considered the ways in which reconstructions of previous incidents were used by some utility companies to drive simulation-based training:

> "The UGI's emergency plan requires each employee who is responsible for responding to emergencies to participate in annual simulation board exercises. Each exercise is prepared by the UGI's distribution engineering personnel and includes scenarios about a system shutdown or loss of a major gas supply line, a shutdown or loss of a district regulator station, or a major line break within the distribution network. The scenario may be based on previous incidents or on incidents described in Safety Board reports. Each exercise must include a step-by-step analysis of the procedures for investigating, pinpointing, and repairing leaks and of the procedures for taking emergency actions and protecting people and property." [589]

These simulations can be used to determine how well teams can cope with the situations that previously confronted their colleagues. However, crews seldom intervene in exactly the same manner as their colleagues. As a result, their actions help to shape new scenarios that differ from the events that occurred in the original incident. Simulation based training, therefore, enables crews to explore more general forms of failure that are based on the particulars of a singular incident. This close relationship between training and reconstruction is emphasised in the Rand report into the NTSB:

> "The NTSB should review its internal technical capabilities to support future accident investigations, including the potential for crash reconstruction and the requirements for system testing in support of complex accident investigations. The safety boards long-term requirements for facilities should include consideration of their use for staff training, recognizing that facilities can serve a dual function." [482]

It is to be hoped that future incidents will not occur in exactly the same way as previous incidents. If they do then this clearly indicates the failure of a reporting system to address the underlying causes of any failure. However, it is not clear how the particular details of an adverse occurrence can be used to anticipate other, more general forms of future failure. The following list identifies a number of ways in which reconstructions can be manipulated to support this form of analysis. The intention is to manipulate the reconstruction in order to either identify training scenarios or to ensure that any recommendations address a wide range of potential future failures:

- *transposition of events.* The most obvious way of generating alternative incident scenarios from any narrative of a particular incident is to alter the sequence of particular failures. For example, in the Garyville incident mentioned in previous sections, a reconstruction might simulate the rupture before, during or after the completion of the controller's transportation command on the batch of unleaded diesel. It is important to stress that the undirected transposition of events will not always lead to failure scenarios. It can also lead to scenarios that might seem extremely implausible. For instance, it seems unlikely that the supervisory control and data acquisition (SCADA) system might generate the high-pump pressure alarm before the pipe failure event. The irony is that many engineers and designers have failed to adequately account for those scenarios that were dismissed as implausible before they occurred [65].

- *omission of adverse events.* A further means of generating alternative scenarios is to omit some of the failures that arose during a previous failure. This can simplify the demands that a training scenario may place upon system operators. Additional complexity can be gradually introduced as teams become more skilled in responding to an adverse situation. This exploits the 'training wheel' approach in which supports are gradually removed from operators as their confidence grows [155]. A particular benefit of this approach is that it can be used to prioritize the allocation of resources to improve system defences. For example, the NTSB report into Minnesota explosion hypothesised that "had the gas line in this accident been equipped with an excess flow valve, the valve may have closed after the pipeline ruptured and the explosion may not have occurred" [604]. This assertion can be tested both using laboratory simulations of the gas flow within the system. Operator performance can also be assessed by reconstructing the course of an incident as though this defence had existed. If the crew can consistently respond to correct and mitigate these alternative scenarios then the proposed defences can be shown to offer some protection. If crews cannot mitigate a failure with these defences then they are unlikely to provide sufficient protection.

- *exacerbation of adverse events.* Reconstructions not only provide scenarios that can be used to assess the effectiveness of potential defences, they can also be used to assess the consequences that may ensure if existing defences are compromised. As we have seen, one of the key differences between incidents and accidents is that particular safety features intervene to mitigate the consequences of failure. We have already argued that an important component of any incident investigation is to determine the 'worst plausible outcome' . These scenarios are again critical both in guiding training and in assessing the potential effectiveness of any remedial actions. For example, the following citation illustrates how NTSB investigators often consider the circumstances that might have exacerbated any failure:

  > "In this accident, the speed and extent of the gas release and fire placed all crew-members aboard the dredging vessels in grave danger. Fortunately, despite the early hour, most crewmembers were awake, alert, and able to respond quickly to the emergency. Given the rapid ignition of the natural gas and the extent of the damage to the vessels, had this accident occurred while most of the crew was sleeping, numerous serious injuries or fatalities may have occurred. The Safety Board concludes that in even a slightly more serious accident, Beans emergency procedures, because they did not require that a precise count be kept of the number of personnel on board the companys vessels at all times, would have been inadequate to account for and facilitate the rescue of missing crewmembers, increasing their risk of serious injury or death." [594]

  Compound simulation techniques provide another means of preparing for plausible worst case scenarios. This approach combines elements of one incident with events that occurred during another previous failure. The result is to create hybrid incidents that blend multiple problems identified during previous incidents. This approach is motivated by Reason's plea not to consider failures in isolation [702].

There are a number of further problems that affect the generalised use of simulations to investigate potential failures. In particular, it is difficult to accurately reproduce operator behaviours under 'experimental' conditions. However, such caveats have to be balanced against the benefits that reconstructions provide in generating the 'what if' hypotheses that direct future development.

8.1.3 Resolving Ambiguity

A key benefit of reconstruction is that it helps investigators to identify omissions and inconsistencies in the evidence that they gather about an incident or accident. This can be illustrated by the NTSB report into the St. Cloud pipeline failure:

> "At about 10:50 a.m. on December 11, 1998, while attempting to install a utility pole support anchor in a city sidewalk in St. Cloud, Minnesota, a communications network installation crew struck and ruptured an underground, 1-inch-diameter, high-pressure plastic gas service pipeline, thereby precipitating a natural gas leak. About 39 minutes later, while utility workers and emergency response personnel were taking preliminary precautions and assessing the situation, an explosion occurred. " [604]

This high-level summary is typical of the sparse information that may be available in the immediate aftermath of an incident. Lack of evidence can prevent investigators from building more detailed reconstructions of the events that contributed to a failure. However, it is possible to use such prose descriptions to help target those events that deserve closer scrutiny. One technique is to scrutinize these narratives in order to identify any ambiguities that require further clarification. These ambiguities partly stem from the flexible ways in which investigators can use natural language to support a number of different interpretations based on the same sentence. For example, the previous quotation includes the observations that the pipe was ruptured by the crew 'while attempting to install a utility pole support anchor in a city sidewalk...'. This abstract description could refer to any number of more detailed procedures that the crew could have been performing in order to achieve their goal of installing the utility pole. They might have been drilling, using a sledgehammer to break the sidewalk, using an auger to secure the anchor etc. If the exact operation that was being performed at the moment of the rupture was critical for a more detailed understanding of the course of events, as is likely to be the case, then investigators must gather more detailed evidence about the crews' actions. The following list, therefore, identifies a number of different forms of ambiguity that can occur in natural language reconstructions, or accounts, of safety-critical incidents:

- *ambiguity of time.* The previous account referred to real-time, '10:50am' and 'About 39 minutes later...'. It also used less precise relative timings that are implicit in phrases such as 'while attempting', 'thereby precipitated'. An important strength of such descriptions during the initial stages of investigation is that it is possible to construct models that describe several different real-time orderings for the events that are identified. For example, the phrase 'About 39 minutes later...' describes possible reconstructions in which the explosion occurred at 38, 39 or 40 minutes after the initial rupture. The scope of the interval is only bounded by the readers' interpretation of 'about 39 minutes'. These slightly vague timings can be made more concrete as further evidence is obtained. However, there are also examples where exact timings cannot ever be confirmed. For example, the time-line of events are often incomplete [589]: Alternatively, if the timing information is not considered significant to the overall analysis

Time	Event
6:48 p.m.	The EPAI foreman called the home of the EPAI Vice President.
6:?? p.m.	The foreman instructed his crew to trace the gas line back toward Utica Street to shut off the gas valve.
6:50 p.m.	The EPAI foreman called the UGI emergency telephone number, advising that they definitely hit the gas line and broke it.

Table 8.1: Excerpt from the Incomplete Time-line of a Gas Explosion

investigators may deliberately choose not to expend finite resources in resolving such ambiguity. All of this contrasts sharply with many of the computer based simulations that we shall explore in subsequent sections, these typically require that investigators commit themselves to precise intervals in which events can occur.

- *ambiguity of place.* The previous account only provides a high level view of the events that contributed to the incident. As we have seen, the US Air Force requires that investigators provide maps of the relative movements of aircraft during an incident. The use of terms such as 'in a city sidewalk in St. Cloud' provide an insufficient level of detail for most investigations. Clearly, any secondary investigation would be expected to produce a more detailed survey of the incident. This illustrates the important observation that any reconstruction will, typically, have to exploit a variety of media if it is provide a complete overview of the many different sorts of information that must support any subsequent causal analysis. Increasingly this may include video footage as well as graphical sketched and textual accounts.

- *ambiguity of action.* The previous summary uses natural language to provide a high-level view of the events leading to the incident. As mentioned, this use of prose provides considerable benefits in terms of flexibility and comprehension. It supports multiple interpretations when necessary evidence is not available. Additional details can be introduced as they are gathered. These comments not only apply to the representation of time and place, it also refers to the account of the crews' actions. Phrases such as 'while utility workers and emergency response personnel were taking preliminary precautions' provide few insights into their actions. Again, evidence must be gathered to determine whether or not their precautions had a significant impact upon subsequent events. There are further benefits of ambiguity in the representation of actions. For example, it is possible to indicate that a crew member performed certain tasks without describing the components, or sub-tasks, that this might have involved. This provides significant benefits if, for instance, these components can be understood from the context of the actions. Problems will, of course, arise when other members of the investigation team do not have the necessary domain knowledge to interpret what this task might have involved. It may also cause problems if, in fact, necessary sub-tasks were either omitted, duplicated or interrupted. Such complexities are masked by this ambiguous action description.

- *ambiguity of motivation.* The previous account provides little information about the potential factors that motivated the crew's decision to anchor the utility pole in that particular location. As with the other forms of ambiguity; there are multiple reasons why natural language descriptions avoid spelling out such factors. In the aftermath of an incident, it can be very difficult to gather objective evidence to support explicit interpretations of individual performance. It is also the case that many investigators lack the human factors training to be confident in proposing more explicit models of the cognitive and perceptual factors that influence operator behaviour. Most reconstructions entirely avoid representing or reasoning about the internal cognitive factors that motivate particular actions. Both natural language descriptions and computer-based simulation techniques, typically, therefore, focus on observable actions only.

- *ambiguity of cause.* The previous description is ambiguous about what exactly caused the incident. It might have been caused by the gas service provider failing to document the position of its pipeline. It might also have been caused by mistakes in siting the anchor for the utility pole. Although these both contributed to this singular incident, it is unclear whether either is necessary and sufficient in the general case. Partly as a result of this causal ambiguity, many investigation agencies deliberately separate the process of finding out *what* happened from explaining *why* it occurred [423]. We have, however, argued that these activities are strongly linked. Reconstructions help to validate and guide causal hypotheses. Later sections will argue that it is, therefore, extremely important that tools and techniques be provided to link these two complementary activities. In particular, it is important that causal ambiguities should not be left in a final report so that the reader is left in considerable doubt about the root causes of an incident.

The previous paragraphs have tried to emphasise that there are often good reasons for ambiguity in the initial reconstruction of an incident. For example, temporal ambiguity can occur because there may not be sufficient evidence to determine the exact moment at which an event occurred. Even in the later stages of reconstruction, ambiguity still plays an important role in the communication of information about complex failures. For example, ambiguity of action can help to abstract away from the exact sub-tasks that an operator or system performed if those sub-tasks can be assumed from the surrounding context of the description and those sub-tasks did not play a significant role in the course of an incident. The following paragraphs summarise several of these reasons why ambiguities may remain in reconstructions of the events leading to failure.

As mentioned, there may not be the evidence available to provide definitive information about the specific course of events leading to an incident. The following synopsis illustrates how in some situations it is only possible to gather superficial facts about the course of an incident. This incident involved a relatively small business jet. Without an advanced flight data recorder or detailed information about the pilot's actions it is difficult to reconstruct the detailed events that led to this incident. The aircraft was destroyed and the instrument rated private pilot was fatally injured. Visual meteorological conditions prevailed at the time of the accident. Winds were 170 degrees at 16 knots gusting to 22 knots. No flight plan was on file

> "The vertical and horizontal stabilizers had some skin wrinkling, but little evidence of ground impact. Both propellers displayed forward bending, chordwise dirt streaks and had dug into the ground, burying the spinners. No engine anomalies were found. No control anomalies were found. Fuel was present at the scene, and all tanks were ruptured. Fuel was found in the lines to both engines." [588]

This incident provides an extreme example of the uncertainty and ambiguity that can arise when investigators cannot access some of the sources of evidence mentioned in previous chapters. However, it is also important to stress that similar problems may also arise from the failure of data recorders. This topic was addressed in Chapter 6.4. Ambiguity is also likely to affect the initial stages of reconstruction before all of the available sources of evidence can be retrieved and analysed .

Ambiguity also occurs if there is genuine uncertainty about the events leading to an incident. For example, the following NTSB incident report describes how it may sometimes not be possible to resolve contradictions in witness statements:

> "During the takeoff roll directional control was lost and the aircraft rolled off the left side of the runway. Heavy braking was applied in order to stop short of a fence and the aircraft nosed over inverted. Both occupants were rated pilots. Their statements were contradictory. It was not determined which pilot was manipulating the controls or serving as pilot in charge at the time of the accident." [586]

In other circumstances, it is often possible to build a number of alternative reconstructions that reflect different hypotheses about the events leading to an incident. The apparent contradictions in the evidence can be addressed by constructing several models; each of which assumes that one particular version of events is the correct one. It is then possible to inspect the resulting reconstructions to determine which version of events is the most likely given the balance of evidence. For example, another NTSB incident report describes how a pilot lost control of their aircraft during an acrobatic maneuver [587]. Some witnesses stated that incident occurred when the aircraft was performing an outside loop. Others stated that the failure occurred during an inside loop. Two reconstructions can be developed to reflect each of these possible hypotheses about the sequence of events before this incident.

Ambiguity also arises when investigators cannot be confident in the evidence that they have obtained about the course of an incident. In extreme cases, this can arise when there are only third party statements about what might have happened. For instance, the following incident report relies upon a witness observation of an aircraft that has still not been located:

> "The pilot signed the pilot authorisation form to rent the airplane on December 25, 1994, about 13:25. Before departure both wing fuel tanks were filled at the request of the missing pilot. The time of departure has not been determined and there was no evidence of contact with any FAA ATC facility. A witness reported seeing a low wing airplane

> about 18:00 local 300-500 feet above ground level flying Westbound. He reported that the engine was sputtering when the airplane flew over his house. The missing airplane did not return to the departure airport..." [590]

In this case the narrative description that 'models' the course of events leading to the failure does not explicitly state that the aircraft observed by the witness was the missing Piper. This ambiguity is intentional; it may or may not have been this aircraft. It reflects the lack of certainty about the course of an incident whose causes could not be determined.

Ambiguity can be used to hide the underlying complexity of particular aspects of an incident. This is important if reconstructions are to provide investigators with an overview of an incident. If all of the details of a metallurgical or meteorological analysis were included then there is a danger that individuals might become 'bogged down' in less salient information. As a result, summaries are supported by further references to other documents that can be accessed to obtain additional detail if required. Chapter 13.5 will describe some of the problems that this style of reconstruction can cause for the readers of an incident report. For now it is sufficient to observe that ambiguity often occurs because of the abstraction or filtering process that is used to construct an overview of complex failures. As we shall see, however, it is critical that this process does not have the side-effect of hiding critical information about the course of events that contribute to a failure.

This section has presented a number of reasons why ambiguity can arise in the reconstruction of safety-critical incidents. In other words, we have shown that there are coherent reasons why investigators may simultaneously provide different accounts of the events that contribute to a single failure. However, this ambiguity also creates a number of potential problems. Firstly, there is a danger that any ambiguity in the initial stages of an investigation will not be adequately resolved by the time that a final report is issued. The previous paragraph described an incident in which it was not possible to identify the events that contributed to a aircraft going missing. In such circumstances, ambiguity cannot be adequately resolved and this is explicitly stated in the NTSB report. However, other incident reports are significantly weakened by ambiguities that seem to have been overlooked or ignored by the investigators. For instance, Johnson describes how one maritime incident report fails to describe what crew members were doing in the critical moments before a collision occurred [415]. In consequence many who read the report were left unconvinced about the investigators condemnations of the crews' actions during that interval.

There is also a danger that ambiguity can lead to misunderstanding. The use of ambiguity and abstraction supports several different interpretations of the meaning of a sentence. However, as a result there is a danger that investigators will read more into an account than was intended by the author. Conversely, they may fail to identify the intended meaning of a high-level reconstruction. It can be difficult to determine whether multiple interpretations reflect genuine uncertainty on the part of the writer or whether ambiguity is the result of necessary abstraction from underlying complexity. For example, some of the incident narratives cited in previous paragraphs do not provide information about meteorological conditions. Others omit information about the role of Air Traffic personnel. In the initial stages of an investigation, it can be difficult for the reader to know how to interpret these omissions. It might be assumed that there were no air traffic events contributed to the incident unless they are specifically mentioned. This interpretation need not be correct, for instance, if air traffic logs were still being assembled. Such problems can be minimised by introducing rules that force investigators to explicitly state when certain events did NOT contribute to an incident. For instance, NTSB incident synopses often exploit this approach. However, many regulators have introduced taxonomies that the categorise many different events that might lead to an adverse occurrence [718]. It is clearly impracticable to explicitly state when each of these events does not contribute to an incident.

A final problem is that ambiguity can arise from the medium in which a reconstruction is presented. As we have noted, natural language offers a flexible and expressive medium of communication. However, this power is achieved precisely because it permits ambiguity. Multiple interpretations are simultaneously supported by the use of imprecise language. Ideally this imprecision can be resolved in the final report on an incident by the introduction of additional evidence as it becomes available. However, as we have noted, there are many instances in which imprecision and ambiguity have persisted into the final versions of an incident report. There are further related problems. In particular, there are some properties that are inherently difficult to represent within natural language. For instance, it can be difficult to describe the way in which concurrent events can simultaneously occur across many of the different distributed systems that are involved in complex incidents. If these events are groups according to the systems that generated them then readers get a good idea of what happened to that particular system over time.

However, it can be difficult to gain an overview of what else was occurring throughout the application at any particular interval. Conversely, if a purely temporal sequence is exploited then it may be easier to see what events were happing at each moment in time. However, readers will have to piece together the individual events that occurred within a particular subsystem. A number of techniques can be used to address these limitations. For example, many incident reports contain text-based time-lines. These are constructed using a tabular form that lists the most salient events that contribute to a particular failure. These are recorded in the order in which they are presumed to have occurred during an incident. The investigator then notes down the time at which each event occurred as an entry in the table. Table 8.1.3 illustrates this approach. It provides an overview of the ways in which various events contributed to a derailment [605].

A limitation with this approach is that particular events can become lost amongst the many different items that are recorded in this tabular overview. As a result, some investigators also produce more detailed tabular time-lines that focus on the events that occurred in a particular subsystem or that influenced particular aspects of an incident. For example, Table 8.1.3 focuses on meteorological conditions during the grounding of a tug [603]: The time-line shown in Table 8.1.3 illustrates a similar approach [601]. In this incident, a table is used to chart the timing of an emergency response to a highway incident. There can, however, be considerable overheads involved in ensuring that these multiple time-lines provide a consistent account. It can also be difficult to ensure that any changes in the ordering of a particular time-line, such as those shown in tables 8.1.3 and 8.1.3, are reflected by consistent updates to an overall time-line, such as that shown in Table 8.1.3. These problems are compounded by the difficulty of gaining an accurate overview from many pages of prose descriptions. It can often be difficult to visualise the flow of events that contribute to particular adverse occurrences. What we need, therefore, are tools and techniques that can be used to explicitly capture properties of an incident, such as the temporal ordering of events, that are difficult to reconstruct using prose narratives. It should also be possible to use these descriptive techniques to identify any potential inconsistencies or ambiguities that might exist in the reconstruction. If these stem from a lack of evidence then either further investigations must be initiated or the final report must acknowledge the ambiguity. If inconsistencies are the result of clerical errors in drafting the report then they should be rectified. The following section introduces a range of techniques that have been proposed to satisfy these requirements for the reconstruction of incidents and accidents.

8.2 Types of Simulation

The previous pages have argued reconstruction techniques help to piece together the evidence that is derived from primary and secondary investigations. This, in turn, helps to determine whether necessary evidence is missing or whether there are any contradictions within the evidence that has already been gathered. Reconstruction techniques can also be used in a more subjunctive fashion. By this we mean that analysts can generalise beyond what is known about a particular incident to assess what might have happened under a number of alternative versions of events. As mentioned previously, we distinguish between two different forms of reconstruction. This chapter focuses on the use of computer-based simulations. These are an increasingly popular means of visualising the events that lead to safety-critical incidents. For instance, sketching tools can be used to derive simple story-boards. Alternatively, digital animation systems support more complex, interactive presentations of adverse occurrences. CAD-CAM tools can also be used to build detailed models of the behaviour of physical systems. Virtual reality systems can be used to drive immersive, interactive simulations with varying degrees of 'realism'. A distinguishing feature of all of these approaches is that they rely on computer-based tools to help analysts visualise the course of an incident. In contrast, Chapter 8.3 looks alternative approaches that do not rely so much on the use of computer-based simulations. A

Time	Events before the accident
01:43	BNSF receives flash flood warning (0001) for the Kingman area.
01:57	Track supervisor for Kingman area is notified.
02:24	National Weather Service issues severe thunderstorm and flash flood warning for central Mohave County, effective until 3:30 a.m. Also, before 3 a.m., weather updates (0002 and 0003) are issued to BNSF, including to watch for flash flooding, effective until 4:30 a.m.
03:39	Crew-members of westbound train Q-LACMEM1-08 report to the BNSF train dispatcher that the rain is letting up at Walapai (MP 501.3), and that they saw water in the culverts.
03:56	Train Q-LACMEM1-08 crewmembers at MP 489.7 report to the train dispatcher that there is no water on the ground and only trickles in the ditch.
04:05	The track supervisor begins his special inspection at MP 516.5, moving in an eastward direction.
04:12	Dispatcher tells track supervisor of Q-LACMEM1-08 information.
04:28	Contract weather service issues update 0004 to BNSF, advising to watch for flash flooding, until 6 a.m.
04:30-04:45	Track supervisor reports from Hackberry (MP 509.4) to the BNSF train dispatcher. He does not report high water. He inspects bridge 504.1. He notes water flowing adjacent to and under the bridge. He does not note any unusual track alignment or take exception to either the east- or westbound bridge.
05:07	Dispatcher reports to track supervisor that eastbound Amtrak train 4 is leaving Franconia (west of Kingman).
05:35	Westbound train B-CHCLAC1-05 passes Walapai (MP 501.3). Shortly thereafter, this train crosses the bridge on the north track at MP 504.1. Train crew notices nothing unusual about the bridge on the south track.
05:46	Track supervisor reports from Peach Springs (MP 465.8) to dispatcher. He says he will clear shortly for Amtrak train 4. He does not report any high water.
05:56	Amtrak train 4 derails at bridge 504.1S.

Table 8.2: Textual Time-line Reconstruction of Events Leading to a Derailment

1:57 p.m.:	winds S-SE/ 25 knots, seas 6 to 8 feet
about 3 p.m.:	winds 26 to 36 knots, seas 10 to 12 feet
4:30 p.m.:	seas 25 to 30 feet
5 p.m.:	winds S-SE /40 to 50 knots, seas 20 to 30 feet

Table 8.3: Textual Time-line of Meteorological Events in a Grounding

Time	Time from initial notification	Action
05:53	00:00	Initial 911 call received by WCSD dispatch
05:54	00:01	EMS dispatched
05:56	00:03	Two Slinger Police Department units arrived on scene
06:02	00:09	SFD command vehicle arrived on scene
06:07	00:14	First EMS unit arrived
06:19	00:26	Flight for Life dispatched from Milwaukee
06:38	00:45	Flight for Life arrived on scene
06:52	00:59	Ambulance delivered first van victim to area hospital
08:01	02:08	Flight for Life helicopter delivered second van victim to trauma center
12:28	06:35	Northbound lanes of US 41 reopened
14:03	08:10	Southbound lanes of US 41 reopened; area cleared

Table 8.4: Textual Time-line of Emergency Response to Road Accident

range of formal and semi-formal notations are used to model the events leading to a gas pipeline explosion. A number of further features distinguish these graphical and textual notations from the approaches in this chapter. In particular, the following computer-based simulations typically lack any formal underpinning. They are the product of iterative development and the subjective introspection of analysts. As we shall see, this provides great flexibility in the range of models that can be constructed. However, it also creates a number of practical problems. For example, in many virtual reality simulations it is entirely possible to break the rules of time and space. The same individual can be represented in two different places at the same time during an incident. In contrast, formal and semi-formal models are often supported by precise rules about what can, and what cannot, be represented. They provide mechanisms that help to identify omissions and inconsistencies, such as that mentioned above.

Some formal and semi-formal modelling techniques are supported by computer-based tools. Computer-based simulations can be directly derived from some abstract notations [720]. The distinctions between computer-based simulations and abstract models are, therefore, often blurred. These distinctions are retained, however, because they help to identify two different sets of concerns about the reconstruction of safety-critical incidents. For instance, formal mathematically-based notations can be difficult to interpret by non-mathematicians. In contrast, interactive simulations can often lead to unwarranted interpretations if analysts are seduced by particular animation techniques.

8.2.1 Declarative Simulations

One class of computer-based reconstructions can be described as ‘declarative simulations’. Declarative models describe aspects of a system that do not change during the course of an incident. At first sight, this definition seems to go the general idea that a reconstruction should provide an overview of the *events* leading to an incident. As we shall see, however, these declarative simulations can be used to illustrate the state of a system before and after an event. For instance, they can be used to illustrate the impact of a component failure upon the integrity of a hardware assembly. Sequences of these more static simulations can, therefore, be used to build up an impression of change over time. It is also possible to integrate declarative simulations with other forms of reconstruction to combine a static model of the system with more dynamic views of an incident.

Maps and Plans

Maps and plans provide important information about the environment in which many incidents take place. They can be annotated to denote the position of key objects and individuals before an incident occurs. Further annotations can be used to indicate the changing position of those objects at various moments during an incident. As with many of the techniques described in this chapter, it is perfectly possible to exploit this approach manually. We have, however, chosen to focus on computer-based techniques because they represent a significant area of innovation and development in incident reporting.

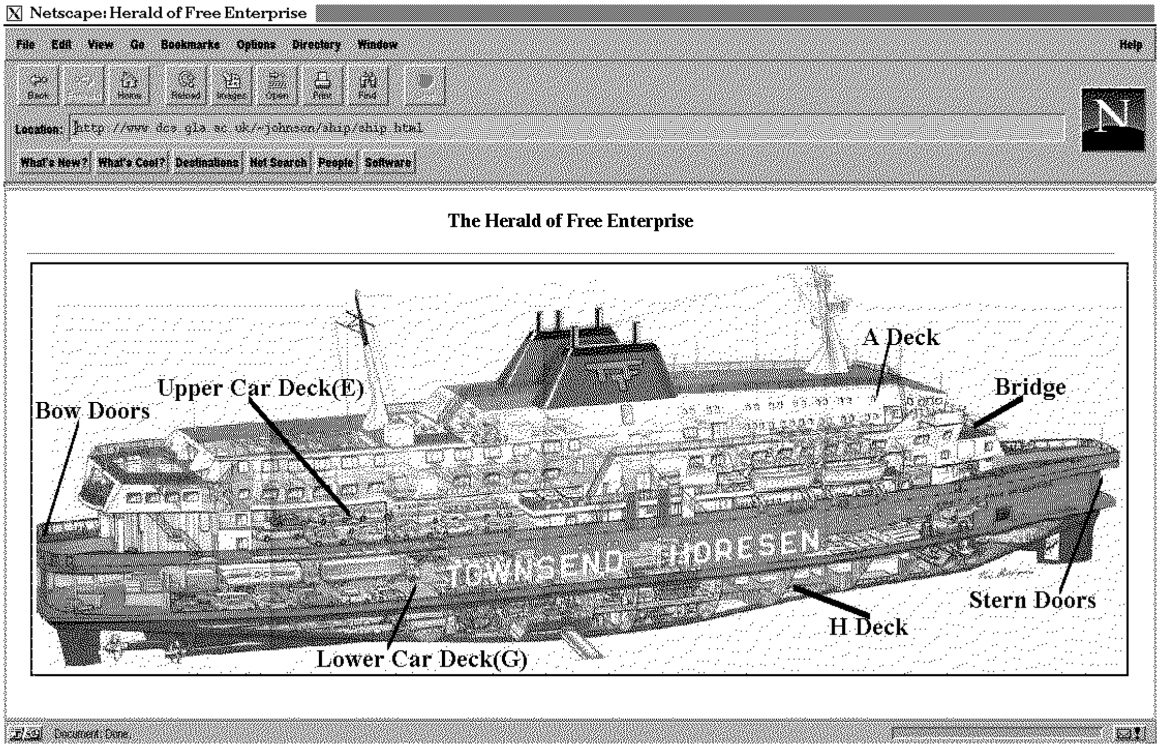

Figure 8.1: Imagemap Overview of the Herald of Free Enterprise

Figure 8.1 illustrates a plan based approach to incident and accident reconstruction. It presents an imagemap of the Herald of Free Enterprise . Thanks are due to The Motor Ship and V. Berris (FSIAD) provided the sectional diagram of the Spirit of Free Enterprise. The interactive reconstruction was developed in collaboration with Anthony McGill [531]. The cross-sectional diagram provides an overview of the layout of the ship that would not have been possible from an external photograph. The labels are used to indicate areas of the ship that played a particularly significant role in the course of the accident. The image is presented in a web browser so that it is available to other investigators over an intranet. This technology provides additional advantages. For instance, users can select areas of the image to request more detailed information about particular areas of the ship.

Figure 8.2 shows how users can select particular areas of the vessel shown in Figure 8.1 to request more detailed information about the incident. This information can take a number of different forms.

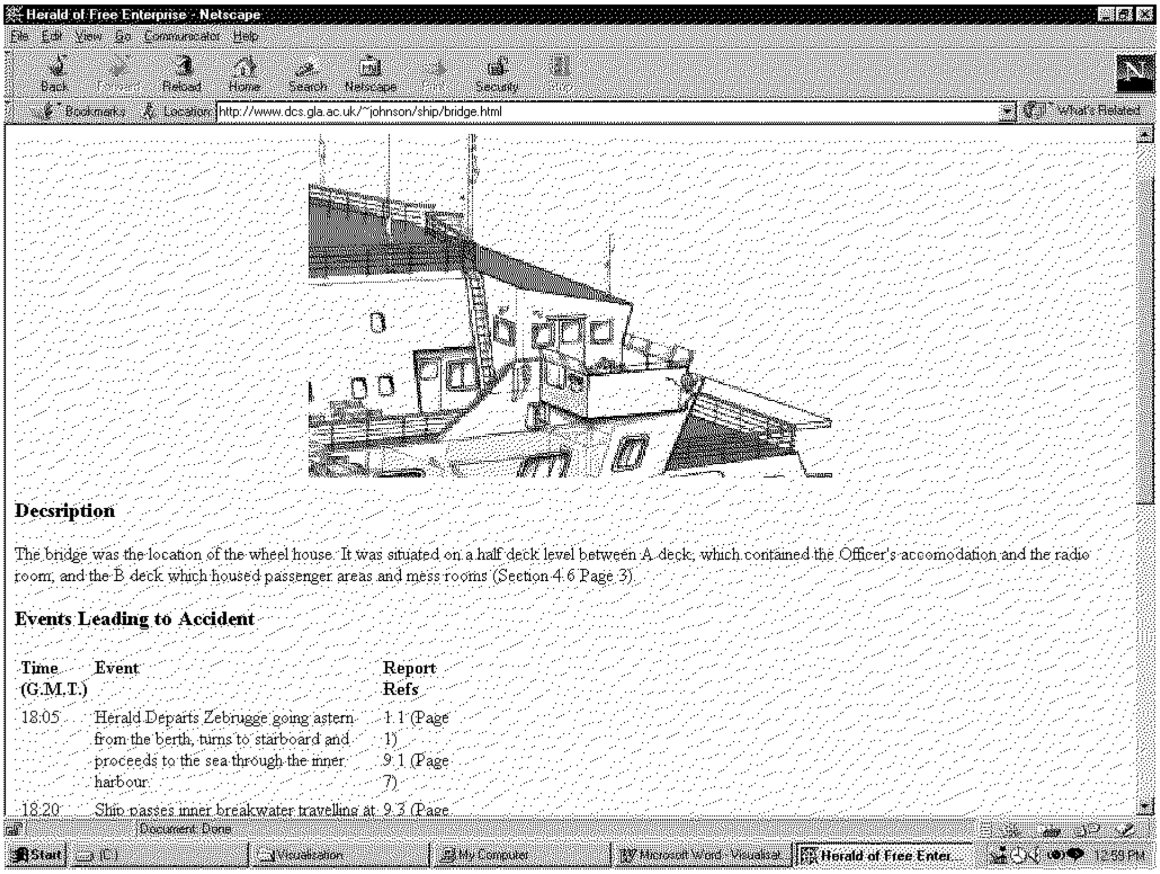

Figure 8.2: Imagemap Detail of the Herald of Free Enterprise

For instance, if the user selects the bow doors they are presented with a range of engineering and construction information:

> "Situated at the bow of the G Deck, the bow doors were double weathertight doors of welded steel construction with a clear opening of 6.0m x 4.9m."

These static descriptions can also be augmented with information about the dynamic events that took place in a particular area of the ship. For example, if the user selects the bridge rather than the bow doors then the screen would be updated to show the textual time-line in Table 8.2.1. This illustrates the manner in which declarative plans, or maps, can be augmented with information about key events during the course of an incident. A time-line can be used to indicate when people and equipment move from one location to another. As can be seen from Table 8.2.1, events can also be annotated to indicate the evidence that supported each observation. This reconstruction was built post hoc and so the citations refer to paragraphs in the Sheen report [737]. However, investigators can exploit the same approach to keep track of the evidence provided by primary sources.

An important strength of the approach illustrated in Figures 8.1 and 8.2 is they avoid the 'god's eye' view that is provided by some computer-based reconstructions. Figure 8.2 only records information that was available to the crew on the bridge. It does not provide information that might have been available to crew on the car deck or in the passenger areas. In contrast, many other simulation techniques provide an overview of all of the evidence that is obtained during primary and secondary investigations. Such reconstructions provide a false impression because they integrate information that could not have been available to any single eye-witness. By using plans and maps to navigate into location-dependent time-lines, it is possible to gain a more accurate impression of

Time (G.M.T.)	Event	Report Refs
18:24	Captain sets combinator 6 on all three engines and the herald accelerates rapidly from 14 knots to possible ultimate speed of 18 knots.	9.3 (Page 7) IV (Page 71)
18:25	Steward hears water on the stairs.	IV (Page 71)
18:28	Ship lurches 30 degrees to port, temporarily becomes stable then slowly capsizes to port.	9.3 (Page 7) IV (Page 68) IV (Page 71)
18:28	Bridge clock stops.	IV (Page 71)

Table 8.5: Textual Time-line Integrated into A Declarative Reconstruction.

an individual's view of an incident. This was significant in the Herald of Free Enterprise accident because individuals on the bridge believed that another member of the crew was supervising the closure of the bow doors while he was actually asleep in his bunk. The 'god's eye' view can then be reconstructed by concatenating each of these discrete time-lines into a single sequence of events.

It is important to emphasise that these declarative models are important not simply in documenting the state of a system prior to an incident, they can also play an important role in documenting the consequences of particular failures. This is most apparent in the use of maps and plans to document key features of major accidents. The UK Air Accident Investigation Branch provide an innovative example of this approach, accessible via [9]. They used a computer-aided design system to model the hull of a Boeing 747. Sections of this model were then annotated to denote the locations where investigators found components on the ground around Lockerbie.

There are more complex examples of map based techniques being used during the reconstruction of adverse occurrences. For instance, police investigators must survey the markings that vehicles leave on a carriageway following an incident [442]. Traditionally, this has involved the use of pencil and paper. Increasingly, however, digital equipment is being used to capture the position of those markings in relation to the layout of the road. This information is then directly downloaded into computer-based reconstruction software. A key aspect of this approach is that many systems support *backwards reasoning*. Information about pre-incident events can be deduced from observations about the consequences of an incident. Many systems use skid marks to deduce the speed and trajectory of particular vehicles. There are, however, a number of limitations with the general application of this approach. Forwards and backwards reasoning from declarative reconstructions can introduce uncertainty. For instance, if the skid marks are eroded or obscured by the effects of the weather or by other debris then it may not be possible to have complete confidence in the results of any consequence calculations. At a more basic level, the calculations that are used to calculate speed and velocity from road markings are, typically, governed by confidence levels. As we shall see, this caveat has a number of important consequences. For example, animated reconstructions can be derived from the information that is deduced by map-based surveying systems. These animations present vehicles travelling at a particular simulated velocities that do not reflect the degree of uncertainty in the initial calculations.

There are further problems with the plan or map based approaches described in this section. Previous figures used an image of the Spirit of Free Enterprise. This is the sister-ship to the vessel that was actually involved in the capsize. We were compelled to use this image in our reconstruction because there is no similar image available for the Herald of Free Enterprise. The sister ship was also used by the official court of enquiry. This example illustrates a more general point. In the aftermath of an incident it may not be possible to obtain a detailed plan or map of particular locations. It may be possible to produce a sketch of the probable location of key objects and people. However, it can be difficult to represent the fact that such locations are based around inferences rather than direct evidence about the scene of an incident. Similarly, the process of developing cross-sectional sketches, such as that shown in Figure 8.1 can introduce biases and distortions of perspective. Photorealistic simulations reduce some of these problems.

Photorealistic Models

s mentioned, plans and maps have long been used to help analysts reconstruct the location in which an incident occurred. Artists sketches can, however, provide a poor impression of the environment, objects and individuals who contribute to safety problems. As a result, photographs are often used to supplement maps and plans. There images provide a more direct impression of the location in which an incident occurred. They can be used to provide detailed information about the physical state of process components, whether they are new or worn, whether they are correctly installed or misaligned, whether they were damaged by an incident or whether they remain intact. Photographic images can also provide an impression of particular environmental factors, such as the line of sight between an operators and a warning signal. Such information can be difficult to convey using plans and maps. It is important to emphasise that such techniques are not immune from some of the biases that sketched images. Different camera angles, exposures and processing techniques can give false impressions about what could or could not be observed during the course of an incident.

As with plan-based simulations, there has been a recent revolution in the use of photographical images to support incident reconstruction. A range of computer-based techniques are creating both opportunities and challenges for investigators. One of the biggest benefits of recent developments is that investigators can take images in the field using digital cameras. These can then be sent from a laptop PC using a modem and wireless telephony to colleagues in other regions country and throughout the globe. in some recent cases this has been done interactively with the field investigator being guided remotely to take live images of the incident site. The resulting photographs can then be archived on servers that can then be accessed by other investigators when they are needed. These images can also be used in subsequent litigation and in any subsequent reports.

As mentioned, these benefits also bring a certain number of concerns. It is possible to falsify conventional photographic images but digital editing techniques make this far easier. A vast range of 'post-production' effects can be achieved with relatively little training. In consequence, many countries have strict rules about the ways in which digital resources can be used both during an incident investigation and during any subsequent litigation. For example, investigators can be required to testify that digital resources have been protected from unauthorised 'tampering'. Digital watermark techniques provide one means of achieving this protection. These watermarks are implemented using an identification code that is permanently embedded into electronic data. The code carries information about copyright protection and data authentication. If the resource is edited in any way then the watermark will be destroyed. Evidence of 'tampering' is then apparent if users cannot extract the original watermark from the electronic resource. The interested reader is directed to Hartung and Kutter's overview of multimedia protection techniques for more information about this and similar approaches [312]. In contrast, the remainder of this section focuses on the use of photorealistic pseudo-3D techniques for incident reconstruction.

One of the problems with static images is that it can be difficult to gain an impression of three dimensional space. This is a limitation of both conventional and digital images. Analysts cannot use these images to 'walk around' the scene of an incident. Instead, they are restricted to the perspectives chosen by the person taking the photograph. This becomes an issue because it can be difficult for investigators to predict all of the possible perspectives that might be relevant in the aftermath of an incident. Even if they expend vast amounts of time and film, it can still be difficult for other investigators to piece together the layout of an environment from dozens of static images and plans. Software engineers have responded to similar concerns within other application areas by developing photorealistic tools for desktop virtual reality (desktopVR). These tools avoid the use of cumbersome helmets and gloves that have been recruited by immersive virtual reality systems. Instead, they attempt to provide an impression of movement in 3D space using conventional input and output devices; keyboards, mice and standard computer displays. Although we focus on this non-immersive approach to virtual reality, many of the comments also apply to incident reconstruction systems that expect their users to wear helmets, gloves and other more complex apparatus. In contrast, QuicktimeVR constructs an interactive simulation from a number digital images and presents them

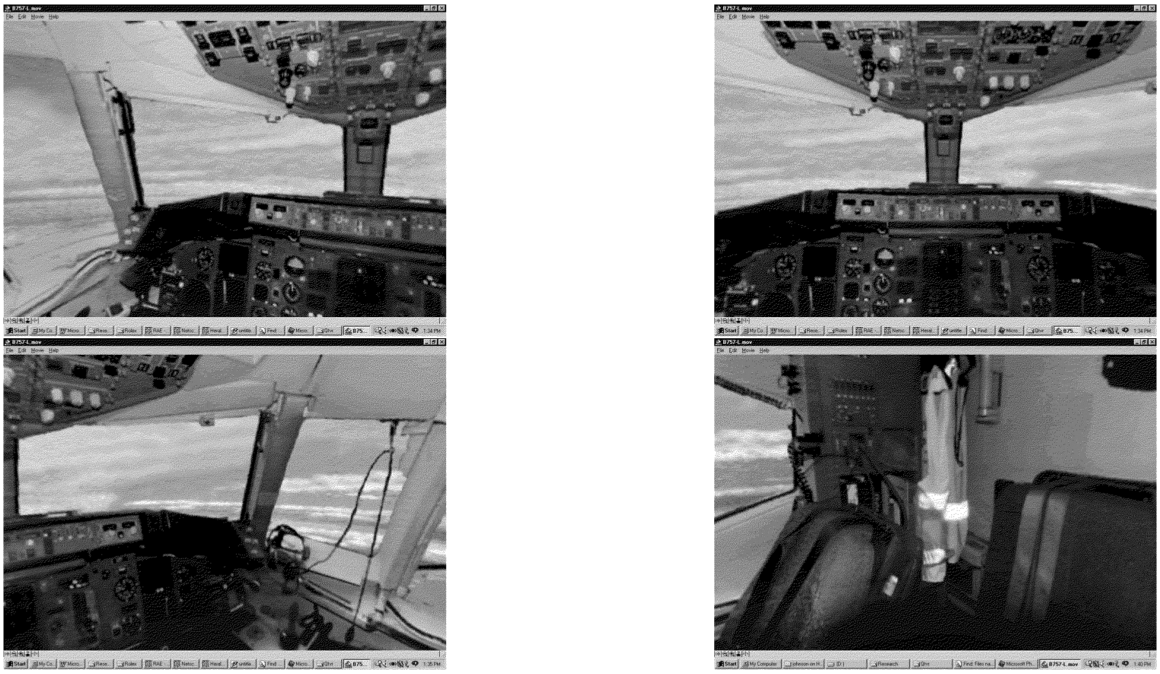

Figure 8.3: QuicktimeVR Simulation of a Boeing 757

on standard desktop displays. These photographs are taken using a motorized tripod which ensures that a still image is taken approximately every N degrees. The value of N is determined by a number of factors including the type of lens that is being used as well as the distance from the camera to the visual horizon. These photographs are then 'stitched' together by the interpolation software. The net effect is to enable users to pan through 360 degrees simply by holding their mouse over the image. Digital effects can also be introduced so that users can zoom in to view particular details of each image. The idea of motion is provided by moving the tripod to a different location. The process described above is then repeated so that the user can again pan around to view the environment from the new location. Figure 8.3 presents images taken from a QuicktimeVR simulation of a Boeing aircraft [425]. The QuicktimeVR images in this section have been reproduced with kind permission of Strathclyde Regional Fire Brigade and were produced in collaboration with Bryan Mathers, Alan Thompson and Bill West [525]. It should be noted that these images provide an extremely poor impression of what is like to interact with the desktopVR system. The resulting simulations not only help incident investigation. They can also provide lawyers, jurors and engineers with an impression of hazardous environments. Unlike films and videos the interactive nature of this approach also enables users to choose their own path through the scene of an incident.

Figure 8.4: QuicktimeVR Simulation of Lukas Spreaders

Figure 8.3 illustrated how the application of desktopVR techniques can be extended from the domains of computer aided learning and scientific visualisation to support incident reconstruction. The Boeing cockpit illustrated in these images represents one of two particular approaches to the QuicktimeVR technique. In this example, a tripod is moved around taking images that form 360 degree sweeps on the scene of an incident. The same approach can, however, also be used with slight variations to provide interactive simulations of process components. This is illustrated by Figure 8.4. Digital editing techniques have been used to joint together a number of still images from a QuicktimeVR reconstruction. The intention has been to provide an impression of the way in which a user can manipulate the software to rotate the object through 360 degrees. The Figure shows how this approach has been applied to the Lukas spreaders that Fire Crews use to extract passengers from road traffic accidents. The process used to create these simulations is slightly different than that used to produce the system in Figure 8.3. Rather than rotating a camera to photograph the environment, in this approach the camera is typically held in a constant position while the object is rotated. By raising or lowering the position of the camera, it is possible for the user to rotate the object around both the x and the y-axes. They can zoom in to view both the top and the bottom of the object.

These photorealistic simulation techniques are declarative because they provide a pseudo-3D impression of the state of the environment or of critical objects at a particular instant in time. However, it is possible to construct multiple simulations to show the state of an object or environment both before and after key events. We have used this most frequently to show the effects of damage or wear on process components. It is also important to emphasise that these resources can easily be integrated into other electronic resources. They can be 'marked-up' in the same way that we annotated the sketches of the Herald of Free Enterprise. This enables users to electronically access linked pages of textual information by selecting areas of the images. Investigators can use this approach to access to the text-based time-lines illustrated in the previous section. This provides an alternative means of introducing dynamic information about critical events into what would otherwise be a static and declarative approach. There are further benefits. For instance, we have used heavy-duty turntables to represent artifacts that are too heavy to be lifted. This enables users to inspect objects such as aircraft engines. They can literally turn these component 'upside down' to select the best angle from which to view any potential damage. Similarly, we are exploring the use of miniaturised cameras to create pseudo-3D models of devices that cannot normally be directly accessed.

At the time of writing this book, these techniques have not been widely exploited to support incident reconstruction. It seems likely that this will change. We are using techniques that are being developed for mass-market computer aided learning systems rather than complex bespoke techniques for incident reconstruction. As a result, these simulations can be developed using relatively low cost technology. QuicktimeVR simulations can be produced at a fraction of the cost of a traditional sketched plan. The examples illustrated in this section were produced by two undergraduate students working in collaboration with members of Strathclyde Fire Brigade. Members of this service have since gone on to develop their own applications of this technology. It is also important to note that the costs of producing these simulations can be defrayed by their multiple applications. All of the systems illustrated in this section have been integrated into the Fire Service's computer-based training schemes.

8.2.2 Animated Simulations

The previous section has presented a range of techniques that support the reconstruction of safety-critical incidents. Maps and plans provide represent the layout of an environment and can be used to locate objects within it. Photorealistic techniques augment these sketches and provide a richer source of information about the scene of an incident. For instance, they can provide evidence about the location of objects that might otherwise have been overlooked when a map is being sketched. However, both of these approaches provide a declarative snapshot of the context in which an incident occurred. They rely upon secondary techniques, such as textual time-lines, to represent more dynamic information. In contrast, the simulation techniques in this section are specifically intended to help analysts reconstruct the changing events that contribute to safety-critical incidents.

Physical Simulations

Previous sections have mentioned that many incident investigations rely upon full scale reconstructions of adverse occurrences. The scale and sophistication of such reconstructions varies from impromptu demonstrations by investigators in the field to full-scale simulations using highly expensive facilities, such as the US Department of Labour's Mine Simulation Laboratory. This facility provides training for safety inspectors and incident investigators using a 48,000 square foot above-ground simulated simulated coal mine. Similarly, NASA's Langley and Ames research centres operate a number of aircraft simulators that attempt to provide a pilot with enough sensory information to convince the pilot that an actual aircraft is being flown. Sensory cues include realistic out-the-window scene generation with 360 degree field of view. The more advanced facilities also provide motion feedback with accelerations applied to the cockpit to simulate momentum shifts in an aircraft. They also employ special seats and suits that are intended to mimic gravitational pull. Other forms of tactile information is simulated by positive force feedback in the pilot's stick and pedals. Acoustic systems simulate natural sounds, such as the wind, and artifical sounds including realistic engine profiles,

Physical reconstructions often involve a high-degree of risk. Reconstructing failures that almost led to a disaster can lead to disaster. Physical reconstructions also, typically, associated with high costs. The NTSB's Wake Vortex flight tests provide an extreme example of this. These were conducted at the FAA's Technical Center in Atlantic City, New Jersey in September, 1995. These vortices can be thought of as a form of turbulence. They can be created whenever an aircraft passes through a section of airspace:

> "Wake vortex: A counterrotating airmass trailing from an airplanes wing tips. The strength of the vortex is governed by the weight, speed, and shape of the wing of the generating aircraft; the greatest strength occurs when the wings of the generating aircraft are producing the most lift, that is, when the aircraft is heavy, in a clean configuration, and at a slow airspeed. (Also known as wake turbulence.)" [611]

Air traffic controllers must, therefore, follow wake vortex regulations that specify minimum separations between particular types of aircraft. These regulations are intended to prevent the following plane from flying through a vortex created by its predecessor [369]. The NTSB's physical reconstructions used a Boeing 727 to generate a vortex. A Boeing 737 was then deliberately flown into the 727's vortex. These were identified using wing-mounted smoke generators on the lead aircraft. The results of the reconstruction were monitored by on-board sensors and from a T-33 chase plane. The risks of such simulations are obvious; especially considering that they were triggered by hypotheses about the causes of a number of previous incidents [611].

The NTSB's wake vortex studies typify the way in which physical simulations are used to obtain many different forms of data. Information was obtained from videotapes, an enhanced flight data recorder, from Boeings portable airborne digital data system, for a 2-hour cockpit voice recorder on the 737 and from test pilot statements. These data revealed that the 727 wake vortices remained intact as much as 6 to 8 miles behind the wake-generating airplane, and wake strength values ranged from 800 to 1,500 feet per second. The video tapes revealed numerous examples of wake vortices breaking apart; linking up; and moving up, down and sideways. This study is also remarkable for the way in which the NTSB exploited the video recordings. These were used not simply to support the investigators' analytical work. They were also used to provide the public with important insights into the conclusions of the final report. This seems to be an increasing trend as incident and accident investigators become increasingly aware of the need to communicate their findings beyond the regulators of their industry [750].

The principle aim of recording this information is to prove or disprove hypotheses about the causes of an incident. The NTSB study was conducted because wake vortices had been implicated in a number of previous incidents. This reiterates the close links between reconstruction and causal analysis. Physical reconstructions help to prove or disprove causal hypotheses. Unfortunately, however, the results from these studies can be highly ambiguous. The failure to recreate an incident may be due to characteristics of the simulation environment rather than weaknesses in the underlying hypotheses. There may also be problems in instrumenting a physical reconstruction to determine whether or not an incident has actually be recreated. Typically, the components and systems that were involved in an incident are heavily instrumented to provide as mush feedback about the potential causes of a failure as possible. This is critical because of the risk and expense mentioned above. Lechowicz and Hunt provide an example of this instrumentation when they introduce a system to provides in-motion weighing, load distribution analysis plus defect detection and classification at wheel, bogie, wagon and train levels [483]. In order to calibrate the measurements from their system they had to run a range of satisfactory and specially assembled bogies over the track. These assemblies provided a mixture of good wheels, defective wheels and vehicle loads and were run at speeds from 30 up to 130km/h. Such calibration tests must be conducted *before* an incident can even be simulated. If they are not then there is a danger that incorrect conclusions will be drawn from the results that are provided by the instrumentation.

Physical simulations supplement any information that can be derived from data recorders that were running at the time of an incident. This live 'incident' data also helps to validate the information that is obtained from physical simulations. If the information derived from a physical simulation is not consistent with that gathered at the time of the incident then a number of potential problems might be diagnosed. For instance, the physical simulation may not have recreated the conditions that held at the time of the incident. Alternatively, the recording devices that were operating at the time of the incident may have been incorrectly calibrated. The recording devices that were running during the incident may not have been correctly calibrated. However, such inconsistencies are not without value. They can often be used to identify alternative conditions that might lead to similar consequences during an incident.

The data that is derived from physical simulations and from incident data recorders can provide parameters for computer-generated models. This has significant advantages. For example, the costs associated with crashing rail trucks at different speeds, typically, prevents a wide range of physical simulations. However, these studies can be used as data points for computer-based simulations that can be run and re-run without significant additional costs. Later sections will identify the strengths and weaknesses of these techniques. In contrast, the following section focuses on the animated presentations that are frequently derived from these computer-based models. In many cases these models are used without the support of physical reconstructions. The high costs, the inherent risks and the problems of instrumenting a reconstruction often persuade investigators to rely upon solely computer-generated simulations.

Computer-Based Animations

We have shown how desktopVR software can 'stitch' together photographic images in order to provide interactive tours of particular environments or rotational images of process components. This is not the only way in which desktopVR technology might support incident reconstruction. In particular, model-based approaches can be used to reconstruct environments that cannot be directly photographed. This is often the case when an incident has resulted in significant damage to a workplace or if that environment has become hazardous in the aftermath of an incident. Model-based approach construct complex process components from a number of geometric primitives, such as spheres and cubes. Application software then renders the image of these composite objects onto the user's screen. The image is updated as the investigator moves around in the virtual environment. In particular, the rendering software must cope with changes in perspective and the occlusion that occurs when one object obscures the image of another.

Figure 8.5 illustrates how this model-based approach can be used to support incident reconstruction. It shows how geometric primitives can be combined to reconstruct complex structures including vehicles and people as well as buildings. This Virtual Reality Markup Language (VRML) simulation was developed in collaboration with Marcus Kramer [461]. This application enable analysts to place construction equipment within a number of different layouts that were proposed for a building site. The equipment could be moved around within these layouts to show users the potential hazards that were posed, for example by overhead cables.

Many model-based reconstructions, such as the one shown above, are developed using declarative environments. They enable users to reconstruct the topography of environment. Libraries can also

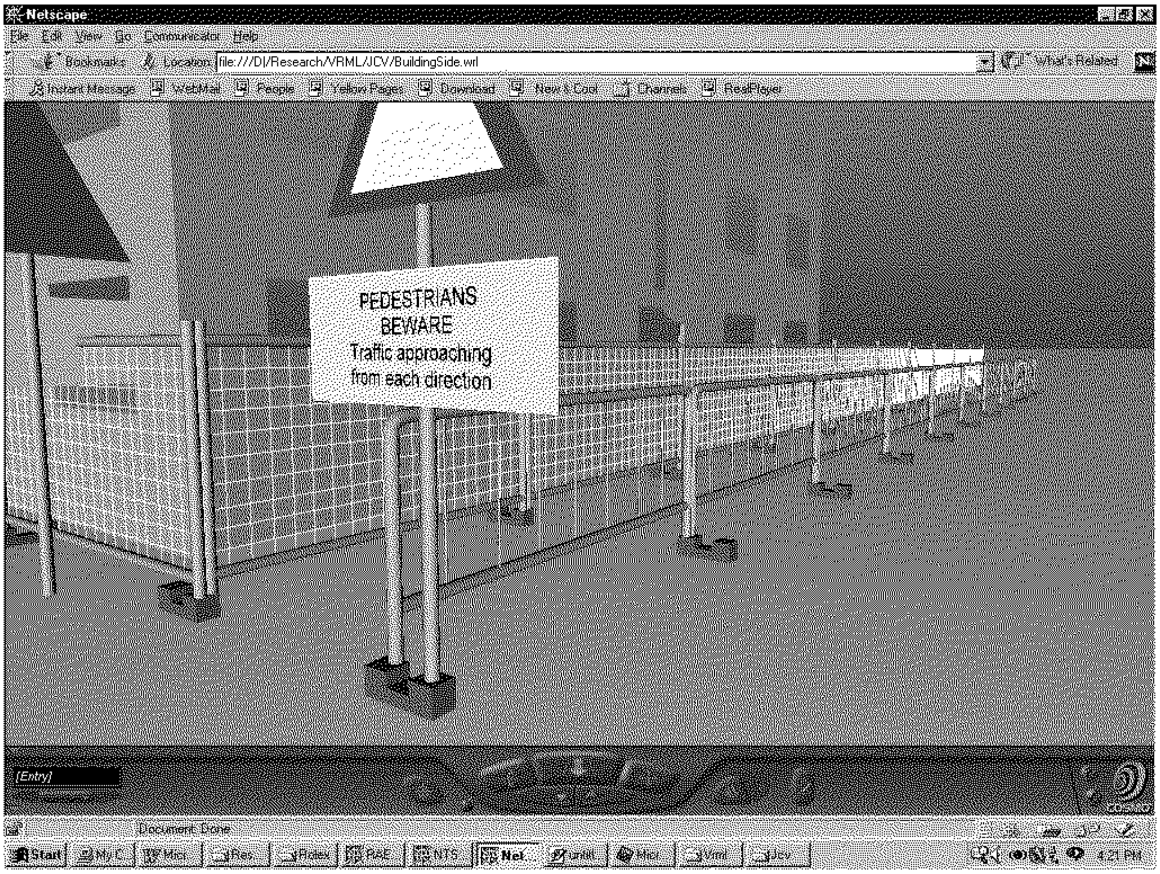

Figure 8.5: VRML Simulation of Building Site Incidents

be constructed to provide easy access to a range of common process components that have previously been developed using the geometric primitives. Analysts must use scripting techniques if these static models are to recreate the 'real-time' behaviour of physical objects. Typically, these programming languages are event driven. Programmers specify the actions that are to be taken when the state of the model changes. For instance, it might be specified that the walls of a building in Figure 8.5 are deformed when a piece of construction equipment hits is driven into it. These programs are developed into simulations in an iterative manner. Any problems in the reconstruction are detected when the program is run and the simulation is viewed. The program can then be amended before being the model is viewed again. Ultimately, however, investigators derive digital animations which simulate the probable course of events leading to an incident. Figure 8.6 presents a number of still images from an NTSB reconstruction that used this approach [612]. This reconstruction had a dual purpose. It was used to validate the inspectors' view of the incident. It was also used to communicate their view of the incident at a public hearing that was convened by the NTSB following the investigation.

A number of concerns affect the production and use of digitised videos from model-based incident simulations. The examples in this section have been produced using extremely flexible modelling software. It is possible to position the 'camera' that determines the user's view of an incident at almost any point in three-dimensional space. This flexibility can result in considerable usability problems for incident investigators who are unlikely to be experts in the manipulation of these cameras. Small changes to the positioning of the viewpoint can result in users entirely missing key aspects of a simulation. As a result, the examples cited in this section use simulations to produce digitised videos that illustrate particular points about the course of an incident. This implies that

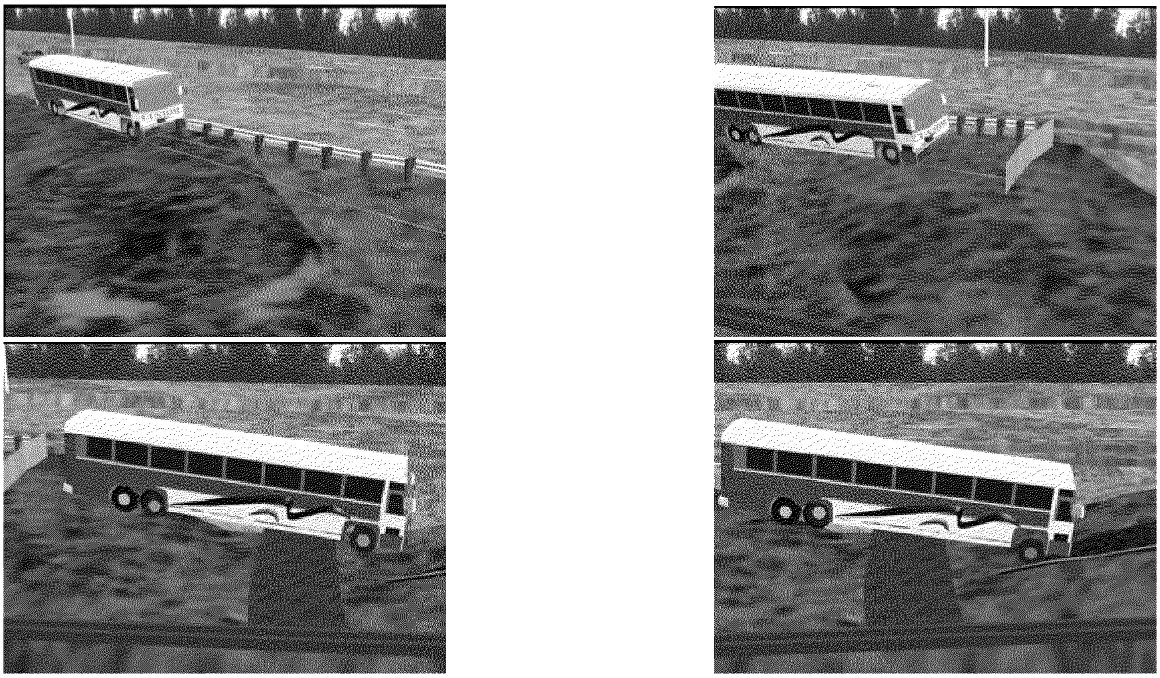

Figure 8.6: NTSB Simulation of the Bus Accident (HWY-99-M-H017).

the pictures obtained from the model are carefully edited to form a linear sequences of images. The end-user cannot change the camera angle nor can then move within a model to obtain new perspectives on an incident. This introduces concerns that investigators may produce videos that are biased towards particular aspects of a simulation. The viewpoint may be place close to the ground to emphasise the apparent speed of a vehicle. Alternatively, the camera may be 'locked' at the same velocity as a moving object so that the viewer gets less of an impression of the relative speed of that object.

Animated models also suffer from many of the concerns that were raised about digitised photographs. It is possible to create a false impression of the events and conditions that contributed to an incident. For instance, model-based reconstructions rely entirely upon the developer to specify lighting parameters. Programmer must ensure that there is a clear correspondence between the world that they *build* and the real situation in which the incident took place. It is for this reason that the NTSB presentation shown in Figure 8.6 is accompanied by the following notice:

> "Disclaimer: Simulations presented below used scene and geological surveys, highway design plans, witness statements, vehicle testing, vehicle plans and vehicle operating characteristics. The depictions represent actual lighting and weather conditions at the time of the accident." [612]

The need to ensure a correspondence between virtual components and physical objects has led investigation agencies and software developers to introduce numerical simulation routines into the scripting languages that drive model-based simulations. The NTSB's models were constructed using tools that were developed by the Engineering Dynamics Corporation. There tools enabled investigators to specify that the bus was initially travelling at sixty miles per hour. The simulation software then reflects the way in which the speed of the bus increased to sixty-two miles per hour when it hit the guardrail. The vault speed of the bus was set at fifty six miles per hour and the impact speed with the opposite side embankment was simulated fifty seven miles per hour.

In spite of the fact that many of these techniques are being applied by the NTSB and other investigation authorities, there are a number of important limitations with this approach. There are considerable costs associated with the time that is required to model even relatively simple objects in model based virtual environments. These costs can be reduced by maintaining libraries of common components. The elements in such libraries may not reflect the subtle differences that exist, for instance between particular models or types of equipment. These costs can dissuade analysts from reconstructing distal causes of adverse occurrences. DesktopVR reconstructions are often biased towards the simulation of those proximal events that have the greatest impact upon their viewer. It can also be extremely difficult to build interactive model-based simulations for certain classes of events. For instance, it is rare to see such simulations that show how managerial or regulatory decisions have influenced the course of an incident. Further problems relate to the use of model-based simulations to represent the outcome of an incident. The NTSB investigators had several good reasons when they decided not to model the impact damage and motion to final rest in Figure 8.6. Firstly, there are computational and technical difficulties in calculating the many different forces that act on complex objects during safety-critical incidents. As a result, simulations must approximate the mechanical and kinematic forces that operate on physical components. This is sufficient for most purposes but can lead to 'unrealistic' effects during impact sequences. More importantly, there are also strong ethical concerns that make the simulation of such consequences unacceptible to those who are involved in an incident.

Abstract Visualisations of Critical Events

Figure 8.7 presents a less familiar application of desktopVR for incident simulation. The interactive reconstruction was developed in collaboration with Anthony McGill [531]. Instead of modelling the objects that are involved in an incident, this approach provides a more abstract overview of the events leading to an adverse occurrence. The image shows three time-lines that recede into the z-plane. Each time-line describes events that are related to a particular aspect of the incident. In this case, systems engineering failures are distinguished from the actions of the chief officer and the assistant bosun. As can be seen, geometric primitives are used to model flags. These are labelled with information about events that affected the course of the incident. Each flag is placed at a point on the time-line. Users can exploit the application software to 'walk' forward along either of the lines. By looking horizontally across the x-plane it is possible to review those events that concurrently affect the other entities that are represented by the parallel time-lines. By looking forwards into the z-plane, the user reviews those events that happen in the immediate future from their position on the line. By turning 180 degrees in the same plane, they can review those events that happened immediately before their current position in time.

As can be seen from Figure 8.7, this technology can be integrated within conventional web browsers. As a result, hypertext links can be associated with the individual objects in the model. Users can select any of the flags to view a page of information about the evidence that relates to that event. This is essentially the same mechanism that was described for the QuicktimeVR visualisations introduced in previous sections. A number of refinements have been proposed for the basic approach described above. For example, we have set some of the flags at a 45 degree angle to the view shown in Figure 8.7. This has been used to provide viewers with a means of identifying when there is conflicting evidence about a particular event. Similarly, some flags are not planted into a time-line. Instead, they 'hover' above the rest of the flags. This has been used to denote events for which there is not accurate timing information.

The approach shown in Figure 8.8 has a number of important benefits. The three-dimensional structures used in the model are extremely simple. Analysts, therefore, avoid the overheads associated with the three-dimensional modelling of complex real-world objects illustrated by Figure 8.5. It is a trivial exercise to add new flags. Investigators simply duplicate the existing components and update its label. There are, however, a number of practical limitations. For instance, it can be a non-trivial exercise to use two-dimensional input devices, such as the conventional keyboard and mouse, to navigate three-dimensional space. These difficulties can be eased by the provision of a number of simple interface design techniques. For example, pre-defined viewpoints can be built into the system so that users can move from one point to another by selecting from a menu of options. This avoids the need for investigators to continually check whether any mouse movements will alter their position or orientation in the X,Y or Z planes. In Figure 8.7 this is achieved by a menu

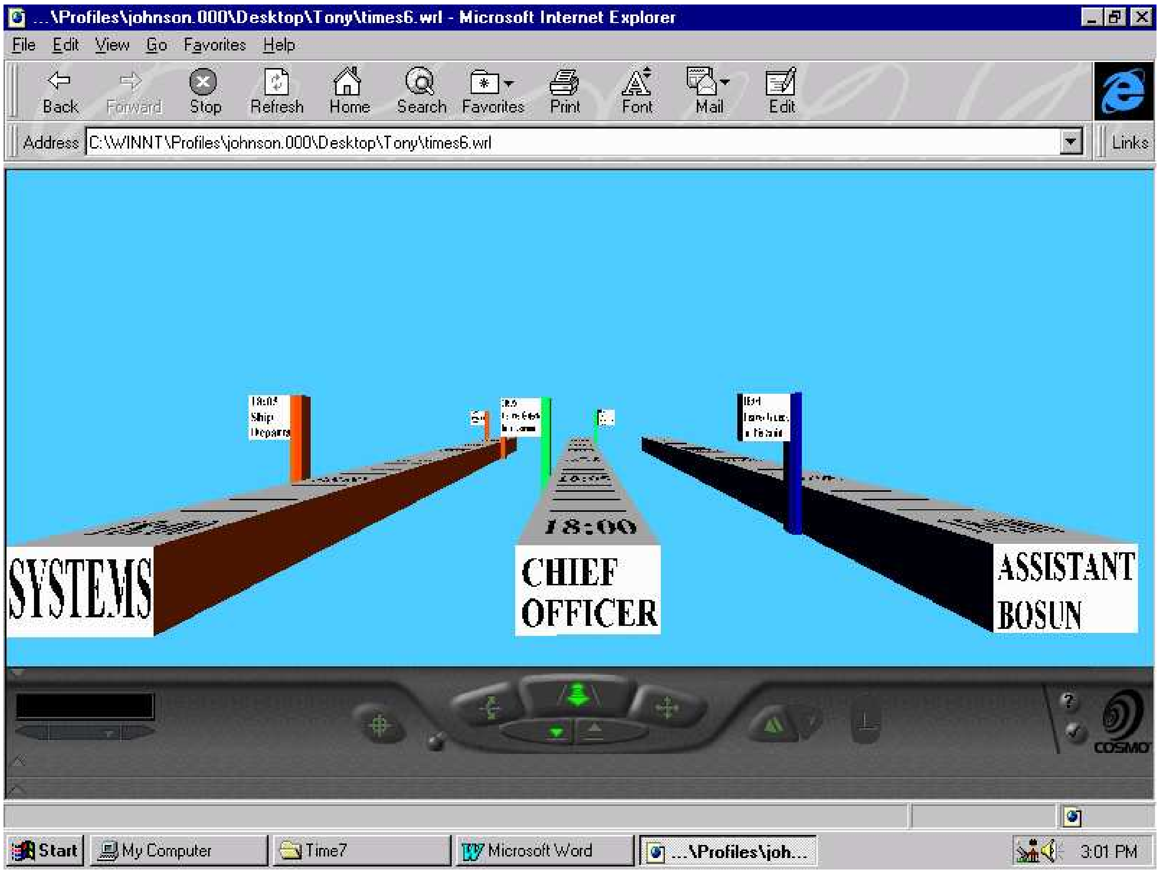

Figure 8.7: 3 Dimensional Time-line Using DesktopVR

of times. If the user selects 18:45 then the system automatically 'walks' them along the time-line to that position. Such techniques can be used to address specific navigation problems. They are not, however, a panacea and it is important to balance the enthusiasm for these new approaches against these practical problems. We cannot expect current generations of incident investigators to instantaneously acquire the three-dimensional skills of computer games enthusiasts.

The techniques illustrated in Figure 8.7 remain the subject of on-going research. Further applications must be developed to demonstrate that they can reconstruct a range of incidents in different industries. Further studies are also needed to determine whether or not there are alternative, more appropriate visualisation techniques. For instance, Figure 8.8 shows how Mackinlay, Robertson and Card's idea of a 'perspective wall' can be applied to support incident reconstruction [509]. This model was developed in collaboration with Ariane Herbulot from the Ecole Superieure en Sciences Informatiques, Sophia Antipolis. The perspective wall makes greater use of the y-plane than the 3-D time-line illustrated in Figure 8.7. However, many of the underlying principles remain the same. Three time-lines are represented on the lower surface of the wall. Events are placed on one of the three parallel lines to denote whether they are associated with a particular subsystem or individual. Their position on the line denotes the moment at which they are assumed to occur. The upper wall is annotated with 'real-time' markers that act as reference points.

Figure 8.9 presents a detail from the perspective wall shown in Figure 8.8. As before, the visualisation is integrated into a web browser. Pages of additional information and evidence can be accessed by selecting individual events. This approach provides a number of benefits over 3-D time-line. As can be seen, concurrent events appear immediately above each other in the x-plane. They can, therefore, be seen at a glance from the position shown in Figure 8.9, which appears in

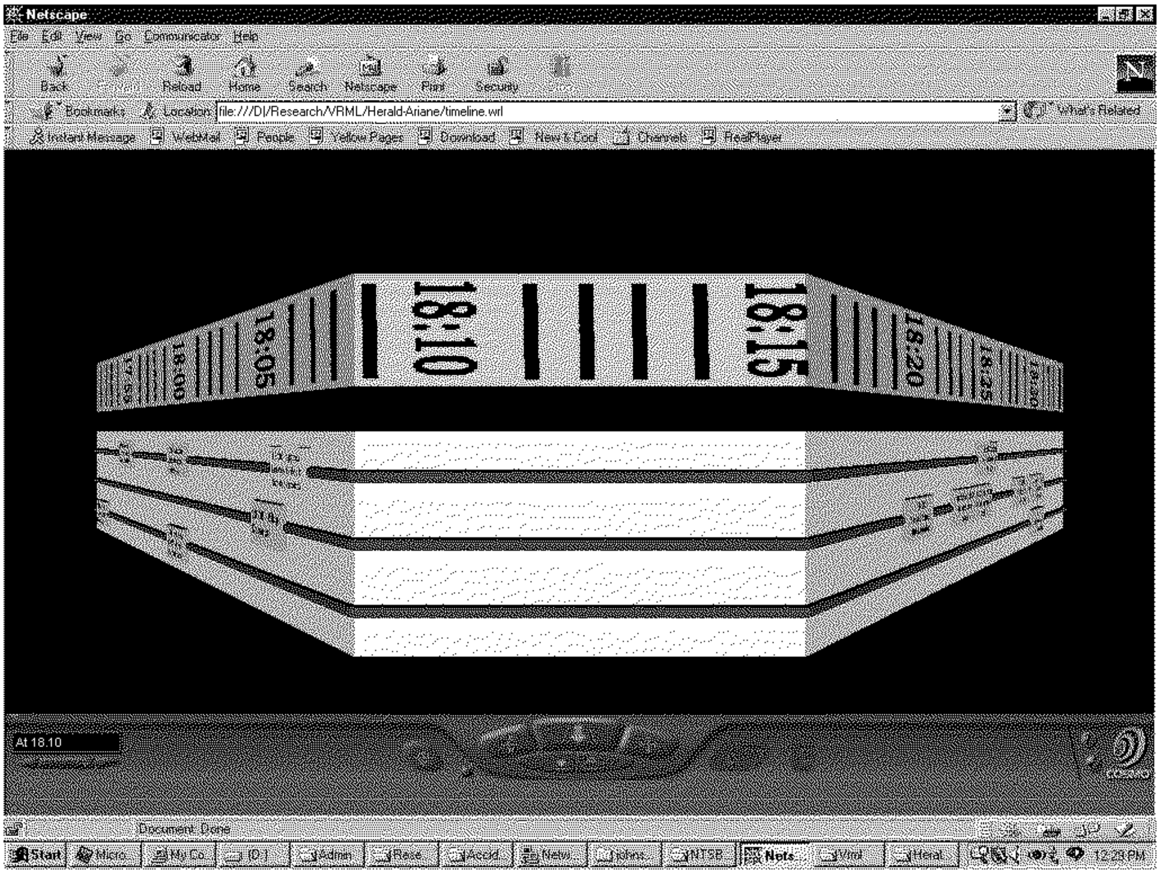

Figure 8.8: Overview of Perspective Wall Using DesktopVR

the menu of viewpoints that was described earlier. In Figure 8.7 concurrent events appeared on the same plance in the z-axis. In consequence, users would have to perform a more complex rotation in order to view these markers. This problem can be avoided if the view point is altered. For example, the analyst can easily see concurrent events on the 3-D time-line if they are positioned directly above the three time-lines. Experience has, however, shown that this can have a disorienting effect on the viewer.

Previous sections have argued that the application of low-cost photographic reconstruction techniques, such as QuicktimeVR, can be extended from mass-market publications to support engineering applications in incident reconstruction. This section has shown how techniques from scientific visualisation and information science, such as the perspective wall, can be applied to similar ends. The meta-level point is that these techniques look beyond current approaches to incident reconstruction. Many of these 'traditional' approaches owe more to the nineteenth century than to the twenty-first. This is a significant and forceful argument because it builds on the underlying analysis of Perrow [677] and Sagan [719]. Chapter summarises their argument that the increasing complexity of many modern systems is leading to increasingly complex failure modes. It is certain that many reconstructions stretch the limits of what can be achieved using manual techniques. The analysis of the Allentown incident, discussed in Chapter 8.3, identified more than 1000 events contributing to human 'error', systems 'failure' and managerial 'weakness'. Some of these events stretched back more than five years before the explosion that triggered the incident investigation. It seems likely that computer-based visualisations will be necessary if investigators are to cope with the burdens imposed by the reconstruction of increasingly complex incidents.

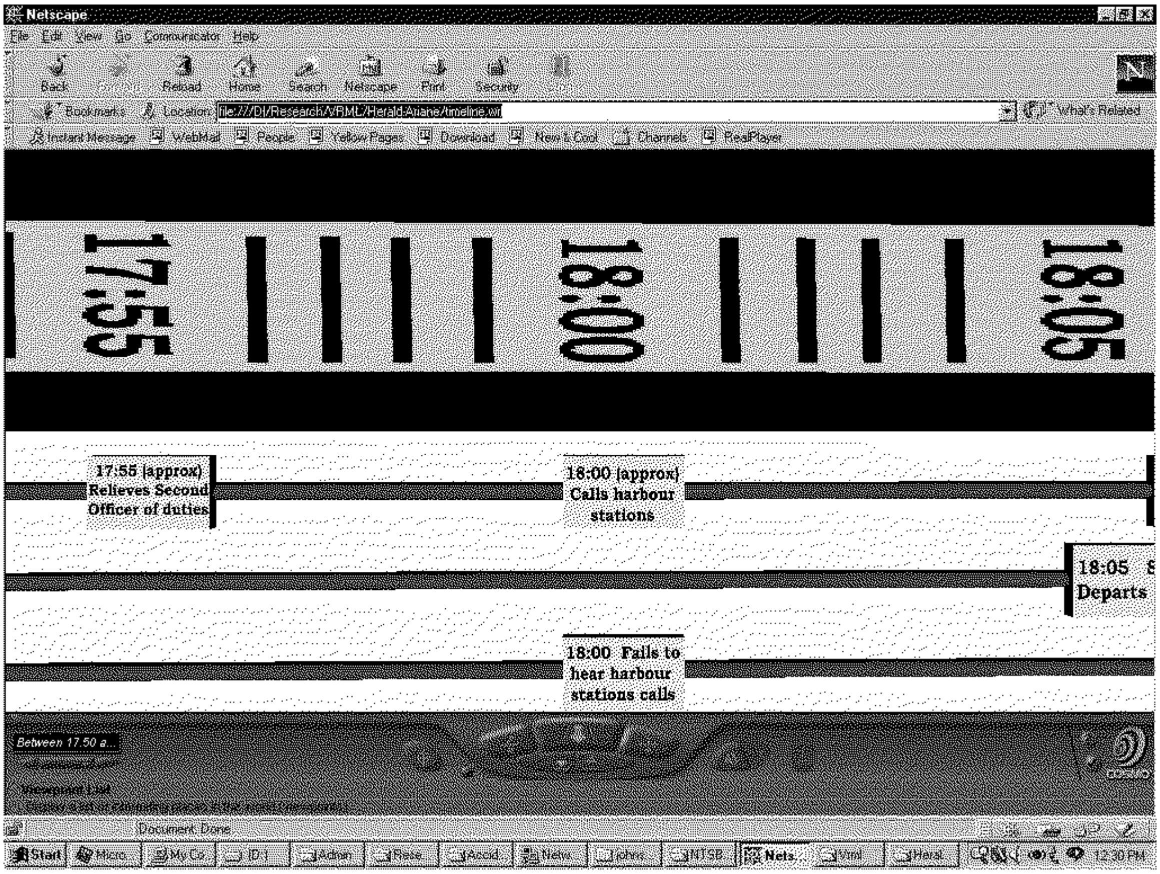

Figure 8.9: Detail of Perspective Wall Using DesktopVR

8.2.3 Subjunctive Simulations

The reconstruction techniques that we have considered up to this point support declarative models of the state of an incident or linear sequences of events leading to an adverse occurrence. In contrast, this section focuses on the use of subjunctive simulation techniques for incident reconstruction. The term 'subjunctive' is used to denote the fact that many simulations are used to explore hypothetical scenarios. They provide a means of analysing what might have happened during an incident. Analysts can adjust the parameters of the model to assess the potential consequences of those changes upon the final outcome of an adverse occurrence. As we shall see, however, the reliability of any conclusions is dependent upon the degree to which the simulation faithfully recreates particular aspects of an incident. It is, typically, impossible to provide completely accurate simulations even using the physical reconstruction techniques mentioned above. In consequence, investigators must analyse which aspects of an incident must be reconstructed if we are to trust the findings of a subjunctive simulation.

Computer-Aided Engineering and Process Based Simulations

Computer-Aided Engineering tools provide one of the most powerful means of deriving interactive reconstructions of complex incidents. This software enables analysts to construct models that are intended to reflect the physical properties of the system itself. For instance, Cole and Cebon have developed a two degree of freedom mathematical model that simulates dynamic tyre forces for tractor-trailor combinations of nine articulated vehicles [172]. The results from this model have been validated in a number of empirical studies and have been used to identify combinations that yield particularly 'strong dynamic interaction'! Others models predict the behaviour of gases within different environments. Numerical techniques derived from computational fluid dynamics have been applied to explain particular combustion processes [551]. An important benefit of this approach is that models, which are used in the design of an application, can also be used to analyse potential incidents. For example, models that explain the properties of fire resistant surfaces, for instance in furnaces, can also be used to explain why other surface fail in similar situations [770]. The dual predictive and analytic nature of such engineering models offer huge benefits, not least in the costs that would otherwise be associated with model development for incident reconstruction.

A recurring theme in this book is that there will be more than one set of events that can lead to the same adverse outcome in complex, technological systems. For this it follows that subjunctive simulations have an important role to play in searching for these alternative paths. Chapter 8.3 will briefly introduce techniques from abstract model checking that have been deliberately developed with this in mind. For now, however, it is sufficient to stress that subjunctive simulations can also build upon, and contribute to, training activities. For example, NASA have developed a simple interactive model of the DF88 directon finder that is used to locate aircraft Emergency Locator Transmitters [567]. Users can interact with this model to explore and practice different location scenarios. The same underlying software can also be used to during lab-based studies that are intended to explore the use of such applications in the aftermath of 'real' incidents. This illustrates the important point that subjunctive reconstructions, in common with the other reconstruction tools in this chapter, should also be capable of modelling the events that occur after an incident has taken place. As we shall see, more lives may be threatened by an inadequate response to an incident than are lost in the immediate aftermath of an adverse event.

A number of limitations affect the use of engineering models to explore alternative scenarios for safety-critical incidents. Many engineering models gain their predictive power by focusing on specific aspects of more complex interactions. Many of the factors that must be considered during the investigation of 'real world' incidents must, therefore, be excluded. For example, computation fluid dynamics can be used to model the combustion of particular materials. These models cannot, however, be applied to predict the impact that the progress of a fire will have upon key electrical subsystems during an incident. Fortunately, a number of companies have developed integrated toolsets that can be used to address such concerns. These tools enable designers to intergrate models from different engineering disciplines. The output of a conflagration model might, therefore, be used to drive predictions about the integrity of electrical subsystems. Figure 8.10 illustrates the interface to Boeing's Easy5 tool. This software can be used to simulate systems containing hydraulic, pneumatic, mechanical, thermal electrical and digital sub-systems. Simulations can be constructed in a number of ways. For instance, functional blocks can be used together with pre-defined components that model physical elements, such as pumps, gears, engines, etc.

Other limitations stem from the implicit assumptions that must be made in orer for these models to be tractable. This is illustrated by the caveats that precede a set of results from simulator studies into aircraft icing incidents:

> "It is important to take in consideration the following assumptions and limitations of this analysis:
>
> 1. The flight conditions just prior to the upset was considered a steady state condition, meaning that all angular rates were considered small and the dynamic aerodynamic derivatives could be considered negligible.
> 2. The Power Effects (Specially the propeller slipstream effect) in the EMB-120 is very strong and for this preliminary analysis was not fully considered when calculating some aerodynamic coefficients.
> 3. The ice effects on the aerodynamic coefficients were taken from wind tunnel test results and only some Reynolds Number corrections were applied.
> 4. The flight simulation (6 DOF) is valid only up to the pusher firing angle of attack (approx. 12.5 deg). Above this angle the aerodynamic data and the effects of any asymmetric flow separation are not valid or not considered.

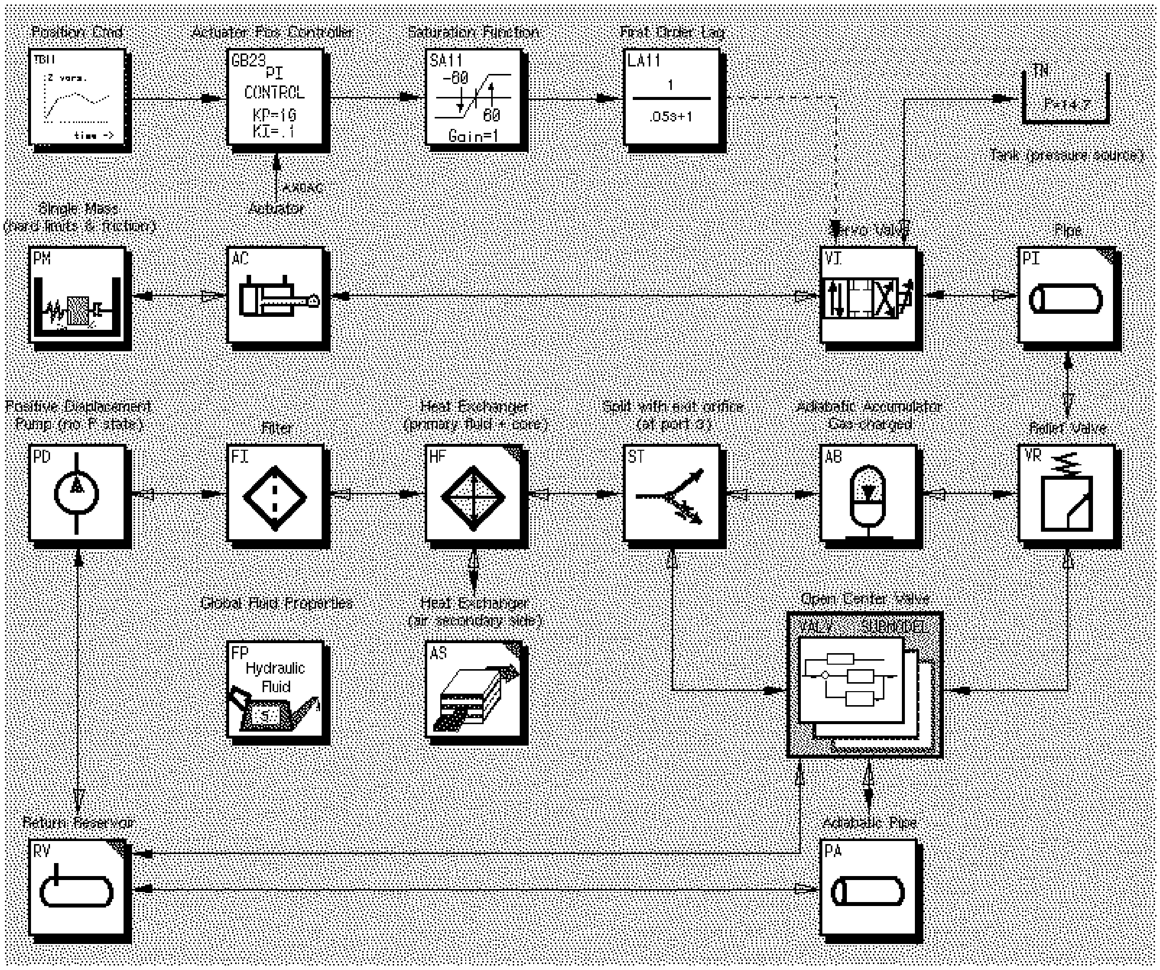

Figure 8.10: Graphical Modelling Using Boeing's EASY5 Tool

5. For this first preliminary flight simulation, only some aerodynamic parameters were modified and for this reason some special assumptions were made due to lack of time.

All assumptions, however, were considered not relevant to this preliminary analysis." (Appendix E, [602]).

There are further limitations with existing systems. Most approaches work best for physical systems that have a linear behaviour. Systems that exhibit nonlinear and time-varying characteristics pose a considerable challenge. Multi-model systems offer a potential solution. These exploit a 'divide and conquer' strategy that represents the behaviour of a complex process in terms of a number of simpler component models. However, there can be considerable instability as a model switches between these component simulations [555]. Engineering models of human operators can be seen as a pathological example of a non-linear system. Subjunctive reconstruction systems have been developed to simulate the potential behaviour of individual operators [440]. However, these tend to be based upon high-level cognitive models rather than the more fine grained control-level simulations of engineering reconstructions. There have been a number of recent attempts to develop hybrid techniques that capture elements both of continuous control in the context of more discrete human decision making, for example in the domain of air traffic control [775]. It remains to be seen whether or not these techniques can be extended from a limited number of case studies to support the reconstruction of more complex incidents and accidents.

In contrast, discrete models of individual and group cognition have been more widely applied to simulate the events leading to safety-critical incidents. They have even been developed to model the behaviour of crowds and entire populations during emergency situations [233]. Much of this work has been inspired by the success of epidemiological modelling that often depends upon assumptions about human behaviour. For instance, Moss has recently extended this epidemiological work to derive simulations of middle management behaviour during critical incident management [552]. This work is particularly significant given that most work on simulation in incident analysis tends to focus on operational failures rather than management or regulatory behaviour. A range of similar models have, however, been developed within the field of management studies to enable users to speculate on the influences that affect key decisions [299]. The theoretical underpinning for many of these multi-agent, predator-prey models is provided by game theory. Individual operators are assumed to pursue independent goals with limited resources under conditions of uncertainty [715]. It is perhaps surprising that few of these models have been used to support incident investigation, especially given the emphasis that Reason and others have placed upon the organisational precursors to failure [702]. On the other hand, it remains to be seen whether such models actually help to analyse different patterns of individual and group performance under high-stress situations. In particular, there continue to be practical and ethical difficulties associated with the validation of these models.

Model Driven Virtual Environments

As mentioned, investigators can use subjunctive simulations to reconstruct alternative hypotheses about the course of an incident. Models of underlying physical processes can be used to replicate the 'real-world' behaviour of complex applications. Evidence that has been gathered during primary and secondary investigations can then be used to parameterise these models. If critical values are unknown then investigators can iteratively inject potential estimates for those numbers. They can then inspect the behaviour of the simulation to validate their estimates. If the modelled behaviour faithfully reconstructs key observations about the course of the incident then the values may be retained. If there are significant differences between the simulation and the observed behaviour then either the values must be revised and the test run again or questions must be asked about the veracity of the observed behaviour.

Bolte, Jackson, Roberts and McComb provide an example of subjunctive reconstruction when they describe the NTSB's use of simulations in road traffic accidents [87]. Data from event recorders can be used as input to a growing range of traffic simulation programs that model the performance of cars, buses, trucks and trains. data These devices are currently only fitted to a minority of commercial vehicles and so the software also relies upon information from witness statements to deduce probable driver inputs. Physical evidence, including the final resting point of the vehicle and any resultant damage, can also be introduced as parameters to the current generation of reconstruction software. All of this information helps to produce crash pulses. These graphs describe the likely acceleration profiles that would be required in order for each vehicle to end up in their final positions with the degree of damage that was recorded. Figure 8.11 illustrates the crash pulse that the NTSB can obtain from road-traffic simulation software.

A collision between a van and a train at Wagner, Oklahoma can be used to illustrate the role that such graphs play during an incident investigation [87]. Witnesses in the van reported that the driver stopped before proceeding over the railway crossing. However, the train engineer contradicted this statement. The train had an event recorder on board and data was obtained from this to determine that it was travelling at approximately 46 miles per hour when the collision occurred. The final position of the van was surveyed in the aftermath of the incident. Data was also collected about the damage that the train had inflicted in the collision. "Reconstructionists" at the NTSB then varied the speed of the van at the point of impact to determine the most accurate trajectory and the related initial speed of the van. This illustrates the way in which subjunctive simulation techniques can be used to explore alternative hypotheses about the events leading to an incident; the van did or did not stop before the collision.

The graphical format used to represent crash pulses, illustrated in Figure 8.11, has a number

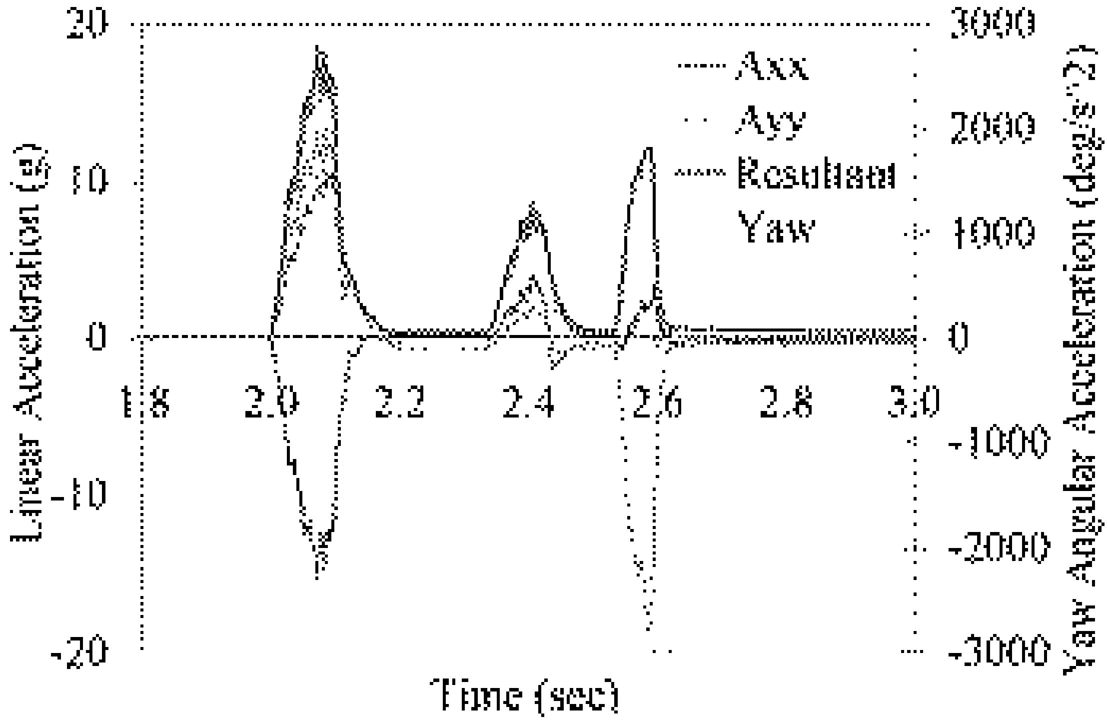

Figure 8.11: NTSB Simulated Crash Pulse Of School Bus and Truck Colliding

of benefits. In particular, it provides a precise numerical profile for vehicle acceleration during an incident. However, it can be difficult to use such representations to communicate the key findings from an incident investigation. This is a particular problem when multi-disciplinary investigation teams must agree about the likely course of events. In consequence, most simulation software provides a range of alternate visualisations. Investigators can use these to show their colleagues what particular calculations imply about the likely course of events leading to an incident. This is illustrated by the model-based virtual reconstruction shown in Figure 8.12. The rendering software that is used to generate this image is very similar to that used in Figure 8.5. However, there are a number of important differences between simple animated simulations and the subjunctive systems illustrated in this section. The former systems rely upon ad hoc scripting techniques to develop a single model of an incident. As mentioned previously, this is often used to derive a linear sequence of images that are edited to provide a movie that can be played during an incident reconstruction. In contrast, subjunctive techniques rely upon physical models of application processes. As a result, it is possible to explore multiple alternative hypotheses about the course of events in an incident by altering the parameters of those physical models.

Simulations, such as Figure 8.11, are not the end-point for inicident reconstruction. Once an accurate simulation has been developed to model vehicle performance, it is then possible to reconstruct occupant kinematics. This is important, especially in road traffic accidents, because it can be used to assess the degree of occupant protection that the vehicle afforded. It is important not to underestimate the importance of such feedback for the design of future safety features.

Operator behaviour must be explicitly considered as a parameter to most subjunctive simulations. This raises many problems. In systems without data recorders, investigators must typically rely upon eye-witness statements. As we have seen with the Wagner incident, these can be contradictory, biased and partial. It is difficult to envisage a set of circumstances in which both the train driver and the van's occupants were correct in their accounts. However, many incidents occur during periods of

Figure 8.12: NTSB Simulation of Motor Vehicle Accident, Wagner Oklahoma

operator inattention or fatigue. In consequence, they may have genuine difficulty in recalling their actions. Most traffic accidents last less than 0.10 seconds. They happen at a speed that makes it difficult for witnesses to accurately comprehend 'vehicle interactions' [87].

The problems of modelling operator interaction in subjunctive simulations are considerably eased if information can be obtained from automated data logging systems. For example, Figures 8.13 and 8.14 illustrate how this data can be used to drive biomechanical simulations of the operator's actions [611, 607]. These simulations represent the rudder pedal positions and leg orientations of the crew. Although this cannot be seen from the still image, the colour of the manikin's leg indicates the force output. For example, blue as used to indicate that no force is being applied to the rudder pedal. Yellow indicates a normal force application while red indicates a larger force than would normally be needed during a flight. Given individual human variability and the sparse data that is often obtained from logging systems in the aftermath of an incident or accident, these models must often be treated with caution. This is evident in the caveats that accompany the NTSB's use of these simulations:

> "These simulations were developed as an educational aid although, whenever possible, the scaling and motions of the manikin and cockpit control were modeled after those of the Boeing 737 event being studied." [607]

Figure 8.14 provides an overview of crew interaction in the cockpit. The NTSB have used annotated similar reconstructions with to include excerpts from cockpit voice recorders as subtitles to the model based simulations. This enables investigators to follow key communications as they watch reconstructs of the crews' interaction with their controls. Such techniques are interesting for many

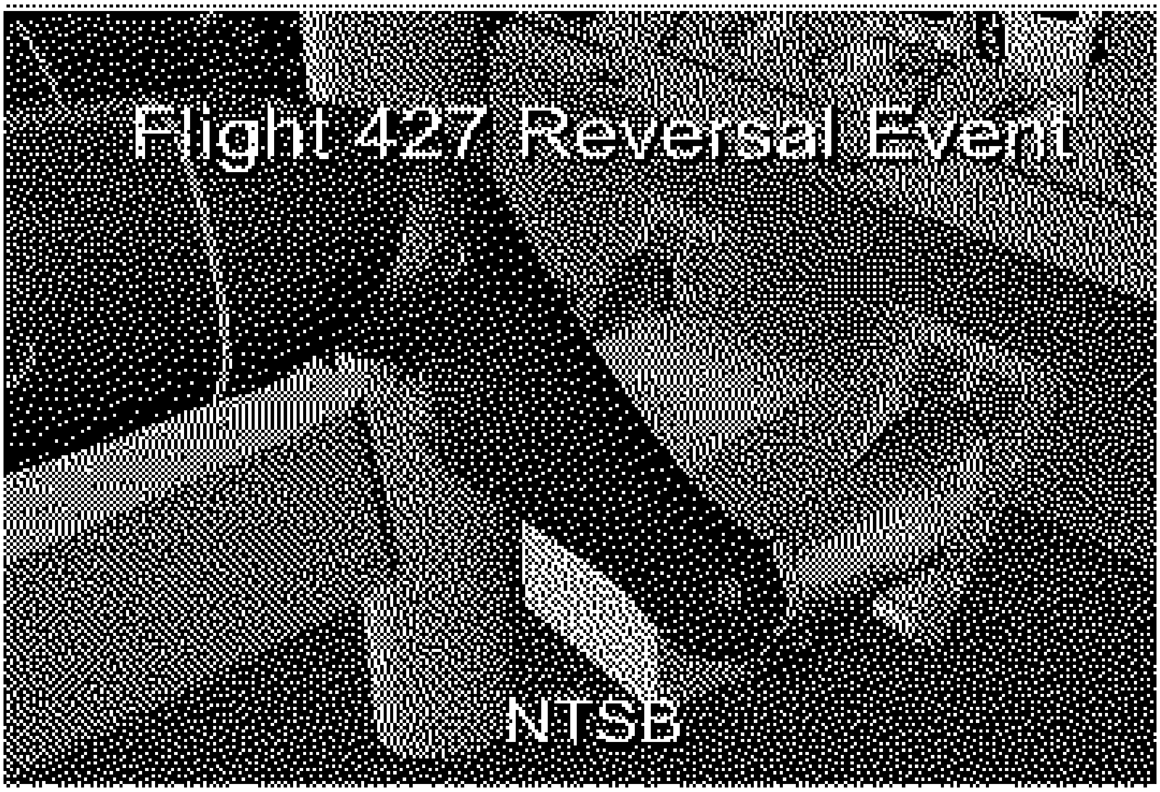

Figure 8.13: Biomechanical Models in NTSB Incident Simulations (1)

reasons. In particular, the public release of these simulations has not created the same controversy as fears about the public release of footage from cockpit video recorders [607]. The US Air Force recognises the distinctions between these different media:

> Animations made from recorder data are not privileged as long as they do not contain (the investigation board's) analysis or input. If the actual audio voices of the mishap crew are incorporated into the animation, simulation or re-enactment videotape, the tape is not releasable due to the privacy interests of the crewmembers or their surviving families." [795]

There are numerous limitations with the subjunctive techniques that are described in this section. In particular, it can be difficult to gain the data that is necessary to derive the models of application components. This means that the resulting visualisations, such as that shown in Figure 8.12, may ultimately rest on little more than guess work. For instance, road traffic simulations typically require information about both drivers' behaviour. This includes steering angles, brake and throttle settings, gear selection, use of engine braking. Some models also expect information about the use of lights, of direction indicators cruise control settings and even wipers or warning horn. In addition, the models require information about the pre-impact speed, engine revs. per minute, acceleration history, braking efficiency activation of anti-lock braking mechanisms etc. If these parameters are not provided then systems either fail to produce a simulation or they exploit default assumptions that may not reflect the conditions that held during a particular incident. Such problems emphasise the point that many subjunctive simulations reconstruct a probable or possible course of events. They do not provide a definitive and unambiguous account of most safety-critical incidents.

Figure 8.14: Biomechanical Models in NTSB Incident Simulations (2)

Probabilistic Simulations

The use of subjunctive simulation techniques is hampered by the problems of data acquisition. Incident recorders are often unavailable. Even if this source of data is available, they often only provide a small subset of the parameters that are required by many reconstructions. The situation is not as problematic as it might first appear. In particular, performance metrics can often be recruited to support data recorders and eye-witness information. Even when these metrics are not available, standard reliability techniques can be used to estimate the potential failure rates for particular pieces of equipment. For example, Table 8.6 presents availability data for the Display Channel Complex computers that formed an important component of the US Air Traffic Management infrastructure. This is calculated as follows:

$$Availability = \frac{Mean\ Time\ Between\ Failure}{Mean\ Time\ Between\ Failure\ +\ Mean\ Time\ to\ Repair}$$

Although Table 8.6 focuses on availability, the same techniques can be used to model more general process behaviour including network loading and resource scheduling. The key point is that the resulting stochastic models mirror the probabilistic behaviour of system components. They can, therefore, be used to determine the behaviour of any simulation in situations where data recorders were unavailable or where their readings provide suspect results.

Monte Carlo techniques provide one relatively simple means of using availability data, such as that shown in Table 8.6 to support subjunctive simulations. This proceeds by generating a random number in the interval between 0 and 1. If the number is less than the availability rate for that component then that component is assumed to be available in the next interval. However, if the number is greater than the availability rate then the component is assumed to fail. For instance, the availability of the Display Channel Complex at the Washington centre in the first quarter of 1994 was 0.9086. If the random number generator produced 0.5 then the simulated Display Channel Complex

Quarter	Chicago	Cleveland	Forth Worth	New York	Washington
Ql, 1992	0.9877	0.8265	0.8899	0.9349	0.9403
Q2, 1992	0.9720	0.9716	0.8059	0.1143	0.9845
Q3, 1992	0.9413	0.9384	0.6984	0.9563	0.9768
Q4, 1992	0.9591	0.9356	0.6566	0.3769	0.9409
Ql, 1993	0.7373	0.9324	0.6628	0.8258	0.8794
Q2, 1993	0.9361	0.9744	0.9294	0.0000	0.9548
Q3, 1993	0.8143	0.9537	0.9011	0.7636	0.9708
Q4, 1993	0.6185	0.9646	0.7397	0.7846	0.9552
Ql, 1994	0.4942	0.9739	0.9458	0.6636	0.9086
Q2, 1994	0.7521	0.8741	0.8257	0.7916	0.8480
Q3, 1994	0.9608	0.9504	0.6955	0.9177	0.9190
Q4, 1994	0.9526	0.9670	0.7996	0.7738	0.9520

Table 8.6: Availability of US ATC Display Channel Complex Computers [591]

would not fail during that quarter. However, if it produced 0.9100 then the simulated component would fail. If the availability increases then it becomes less likely that the random number will be greater than this revised figure and so the component becomes less likely to fail in any simulation. If the availability decreases then it becomes more likely that the random number will be greater than this figure and so the component becomes correspondingly more likely to fail in any simulation.

The previous paragraphs have presented a simplification of the probabilistic techniques that investigators can use to support subjunctive simulations [27]. The key point, however, is that these models help to ensure that reconstructions are based upon observed behaviours even when there may not be direct data about the course of an incident. In particular, information about previous failures can be used to bias simulations towards particular traces of events. For example, if a particular installation had continual problems with their Display Channel Complex then the availability figures would be considerably reduced during reconstructive simulations.

There may seem to be a paradox in the previous discussion. We have argued that subjunctive simulations help investigators to explore alternative hypotheses about the course of an incident. However, this section has argued that reliability data and probabilistic simulations can be used to ensure that the behaviour of a simulation is narrowly based upon the observed behaviour of particular components. This apparent paradox can be resolved by emphasising that these techniques still enable investigators to identify a range of 'what if' scenarios. For example, Monte Carlo techniques inevitably support some of this exploration because different random numbers should be generated on each pass. It should be apparent, however, that rare events will still be very difficult to simulate under this approach. In simple terms, if a component has an associated availability rate of 0.99 for a particular time period then there is 1 in 100 chance that it will fail in any particular run of a simulation. If an investigator wished to focus on those scenarios in which this component were known to fail then they might make a corresponding reduction to the associated availability of that component. The danger here is that such changes might not accurately reflect the availability of that component during an incident. The more changes that are made, the further an investigation moves from any available reliability data. In consequence, it is important that investigators keep a log of the sources of data that are used to drive any simulation together with a detailed justification of any changes that are made to such data.

Probabilistic simulations support subjunctive reconstruction in a number of further ways. One of these approaches can be illustrated by changes in the maintenance procedures for the Display Channel Complex equipment, mentioned above. Approximately five years before the figures in Table 8.6 were published, one of the US ATM sites conducted a hazard analysis for these components. This raises a number of concerns. For example, the standard maintenance practice was to immediately intervene whenever one of the three computing elements failed. The review demonstrated that this might disable the two remaining computing elements and thereby jeopardise service provision. In consequence, a revised plan required that the engineers wait for a low traffic period to begin repairs. In consequence, there was a deliberate decision to continue operating with less-than-full redundancy. Local technicians and managers believed that this is 'less risky' than beginning immediate repairs. Table 8.6 presents availability data for a fully redundant Display Channel Complex. During subjunctive simulation, it would be relatively straightforward for investigators to replace this with information with availability data from sites that operated the revised maintenance policy. Such techniques illustrate one of the ways in which simulations can be used to generalise beyond the specific circumstances of a particular incident to examine other failure scenarios.

These applications of probabilistic simulations must be considered against a number of limitations. Chapter 2.3 has already described the practical and theoretical problems that are associated with the stochastic modelling of software failures. These concerns form part of a wider scepticism about the probabilistic modelling of safety-critical incidents. For instance, Chapter 1.3 has described Wright's recent work in the railway industry [875]. She has shown that previous information about a large number of relatively low criticality failures often provides few hints about the likely state of a system during higher criticality incidents. These caveats have recently persuaded a group of researchers under John Fox at the Imperial Cancer Research Fund to look at *possibilistic* simulations. The intention is not simply to explore what is likely but rather to determine what is feasible given a particular model of the system. The obvious drawback to this approach is that many possible scenarios may lead to a particular incident. Possibilistic simulations, therefore, depend upon arguments about the plausibility of a particular scenario given what is known about a set of events. The emphasis is, therefore, placed upon a more general argument of plausibility rather than a search to quantify reliability estimates. These ideas need to be more fully developed. In particular, more work needs to determine the appropriate form for convincing and plausible arguments about the veracity of a simulation. As we have seen, many reconstruction techniques offer ample opportunities for biasing simulations towards particular viewpoints about the causes of complex failures. There are, however, strong links between this approach and, for instance, the increasing use of model checking in safety and reliability analysis. This approach models an application in terms of a number of possible state transitions. These transitions describe how the state of a system can change over time. Analysts can then specify properties or theorems that they would like to hold over the system. Automated model checking tools will then explore the possible states of the system to determine whether or not the property actually does hold for this model of the system [193].

A number of further issues complicate the application of subjunctive simulation techniques for incident reconstruction. Arguably the most important of these surrounds the modelling of team-based interaction during adverse occurrences. Previous paragraphs have explained how cockpit voice recorders and flight data recorders can be used to integrate human factors observations into physical models of safety-critical incidents. However, we have not explained how reconstructions might be developed when such data is unavailable. Probabilistic techniques, such as Monte Carlo simulations, can be used to address this limitation. System operators can be asked to interact with a simulation as it reproduces a probable failure scenario. Operator behaviour can be monitored over successive runs to identify a range of possible responses. This is, however, a partial solution. Operators often alter their behaviour if they know that they are being monitored in the aftermath of an incident. Such problems are exacerbated if operators must interact with computer-based simulations rather than with their 'real world' counterparts. Moving a desktopVR model of a train around a simulated track is hardly comparable to the physical and mental processes involved in driving a real locomotive. Similarly, it can be extremely difficult to reconstruct the working pressures and team-based interactions that characterise most working environments.

Single and Multi-User Simulations

It can be difficult for individuals to recall their actions immediately before an incident. Even if they can remember, personal, organisational and social pressures can prevent operators from rendering accurate accounts of their behaviour. In consequence, it is important that investigators have some means of reconstructing alternative hypotheses about the role of human intervention in an incident.

There are two main approaches to this form of simulation. The first is to study human interaction with computer based reconstructions, in the manner described above. The second is to derive computer-based simulations that attempt to model human intervention with complex systems.

The following list summarises some of the problems that arise when studying operator interaction with simulated systems. These problems stem partly from the difficulty of staging such reconstructions and partly from interpreting the observations that can be derived from them:

- *problems in obtaining access to appropriate staff.* In the aftermath of an incident, it can often be difficult for investigators to obtain the cooperation of operators who are willing to participate in simulation studies. Such volunteers must possess the necessary skills and experience to take part in the studies. They must agree to have their performance monitored and recorded while they interact with a simulation, even over prolonged periods of time. They must be prepared to 'suspend disbelief' when mock-ups replace application processes. Even if such individuals can be found, it is often important to secure the support of trades unions and other forms of worker representation before many workers will participate in such studies. This usually involves assurances about the ultimate use of any data that is to be derived from the studies;

- *difficulty of simulating individual factors.* The list of requirements in the previous paragraph can have a paradoxical effect on participant selection for simulator studies. They imply that potential 'subjects' must have an active interest and involvement in safety issues. Individuals who are prepared to have their interaction monitored over prolonged periods of time may not provide results that can be generalised across the rest of the workforce. Even if such biases can be counteracted, it seems unlikely that investigators will be able to address the diverse range of behaviours that characterise individual responses to the stress and uncertainty of many incidents. In other words, simulator studies may only provide a partial glimpse of the operator behaviours that might have occurred during a particular failure;

- *difficulties in repeat simulations.* Operator interaction can be observed during several trials with the same scenario. This helps to reduce any nervousness during an initial observation. It can also help to reduce any effects that stem from differences between a simulation and the 'real world' system. However, operators may gradually transfer knowledge gained from previous trials to guide their interaction with subsequent simulations. This not only applies to repeated trials with the same scenario, knowledge can also be transferred between different simulations. It can, therefore, be argued that operator interaction can provide less and less realistic results as the number of trials increases. Eventually there is a danger that boredom and fatigue may provoke behaviours that would not be apparent in other circumstances;

- *difficulty of simulating contextual events.* Chapter 2.3 briefly introduced the range of performance shaping factors that can affect operator behaviour. These can include heat, noise, light-levels, individual fatigue, alcohol and drugs etc. In anonymous reporting systems it is often impossible to determine whether ot not these played a significant role if they are not explicitly mentioned in an incident report. Even if investigators can interview operators, they may not be aware of the extent to which these factors have influenced their actions. Assuming that investigators have identified the importance of heat, noise etc, they then face the task of accurately recreating these influences in a manner that is ethically acceptible. This is less easy than it might at first appear. In some studies, operators exhibit a relatively small number of errors under ideal conditions. In consequence, investigators have started to bombard them with noises, flashing lights and other forms of distraction so that the simulations become parodies of the environments that they represent;

- *difficulty of provoking rare behaviours.* Many of the operator behaviours that exacerbate or cause incidents are extremely rare. In most cases, systems are designed to mitigate the consequences of such actions. Individuals may also be preselected and trained to reduce the likelihood of such intervention. In the aftermath of many incidents, operators can also be sensitised by rumours of their colleagues' intervention. This may make individuals even less likely to repreat those previous failures. All of these factors may force investigators to run many hundreds of trails before evoking a response that might have contributed to an adverse occurrence;

- *difficulty of interpreting observations of simulated behaviour.* Even if it is possible to overcome many of the practical barriers that frustrate the observation of user involvement with incident simulations, there are still many problems associated with interpreting the behaviours that are identified. For example, the same user may interact in one way during one trial but then interact in a quite different manner during a subsequent run. Alternatively, two different operators may react in quite different ways to the same simulation. A range of techniques can be used to examine these differences. For instance, individuals can be asked to explain the motivation for their decisions as they interact with a simulation. However, this process of introspection may force them to re-examine their decisions in a way that would not occur during normal interaction with the system. Alternatively, cognitive modelling techniques can be used to represent some of the psychological processes that might motivate individual interaction [122]. However, it can be difficult to ensure any form of inter-analyst consistency using this approach. There are often considerable disagreements about the underlying mechanisms that provoke particular operator actions during simulator studies;

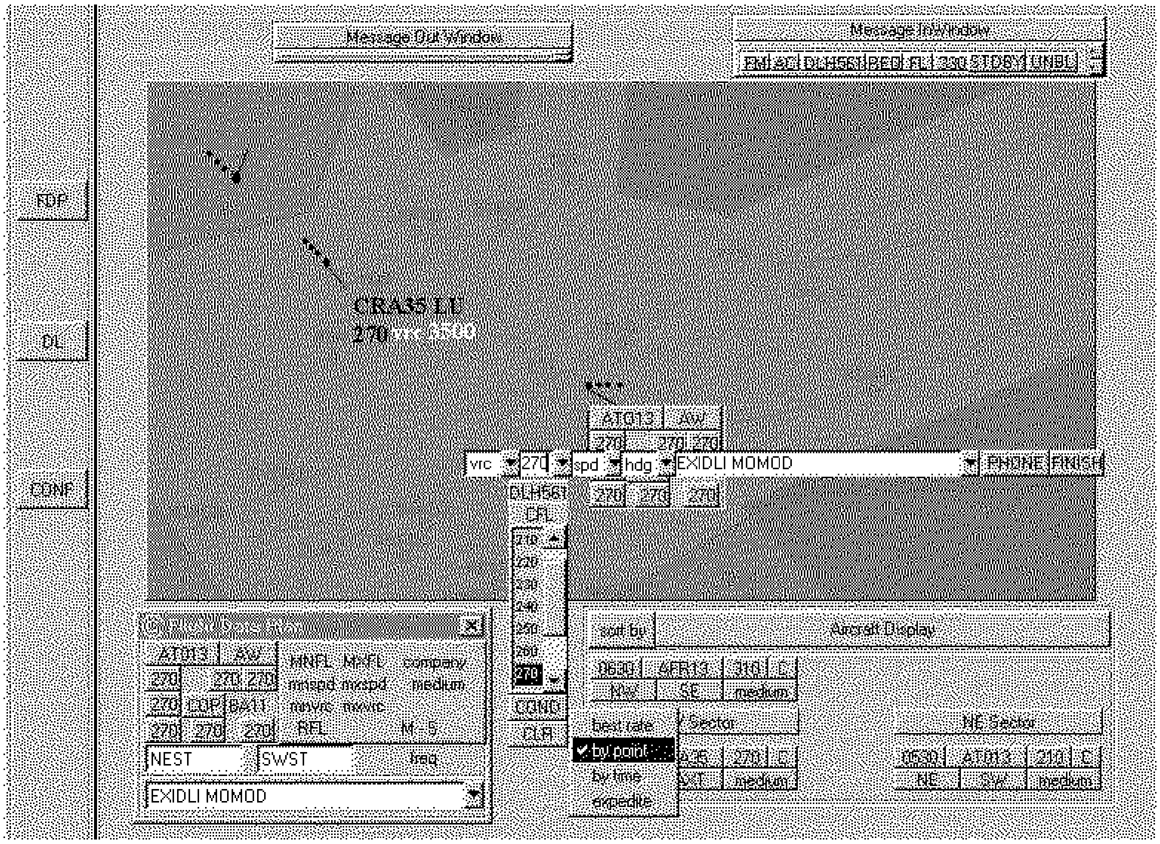

Figure 8.15: Multi-User Air Traffic Control (Datalink) Simulation

- *difficulty of simulating team dynamics.* Previous paragraphs have mentioned the problems of simulating group based interaction. Each of the previous items in the list could be re-written specifically in terms of the difficulties associated with team based simulation. For instance, it is difficult to underestimate how complex it can be to interpret those factors that motivate particular group-based activities. In many incidents this involves a recursive dependency in which my actions can be influenced by which I think that you are thinking about my current intentions [407]. As mentioned in Chapter 2.3 team-based interaction can be dominated by particular individuals. This may or may not reflect what actually happened during a particular incident. It can often be difficult to recreate the balance of skills and personalities that contributed to an incident. Similarly, the problems of securing operator participation are exacerbated by the pragmatic difficulties of ensuring that different operators are all available at the same time. There can also be significant costs associated with the development of multi-user simulations. Figure 8.15 illustrates an interface that has been produced using one of a number of environments that are intended to reduce these costs [720, 721]. This multi-user interface builder has the added benefit that it can be used in conjunction with the model checking tools that were mentioned in the previous section. This example shows an en-route controller's view of a datalink air traffic control system. They can communicate with other controllers and with individuals playing the roles of the aircrews in their sector. The expense involved in running a simulation session with qualified operators should be obvious. The complexity in implementing such a system and then linking it to accurate traffic models should also be apparent.

Given all of these problems, a growing number of researchers have turned to alternative means of simulating operator interaction with safety-critical systems [728]. Rather than creating environments that are intended to enable human operators to replicate the events leading to failure, this approach extends the scope of computer-based simulations. In particular, research teams have developed a number of computer-based simulations that are intended to recreate individual behaviour. Some of these models have also been 'hooked up' to other simulations of system behaviour in order to observe potential interaction. Again these techniques are still in their infancy. For instance, NASA Ames research centre recently reported on the APEX project that was specifically intended to make Human-Machine system simulation a practical engineering technique' [704]. This project has identified an architecture for future systems. The interpretation of auditory and visual input helps to determine the contents of the operator's working memory. This, in turn, influences the process by which an individual selects their actions for intervention. These actions are executed through components that model gaze and attention as well as vocal commands and the operator's gestures. At present, this work focuses on modelling the behaviour of individual operators. It has the important advantage of explicitly representing theories about the factors that motivate operator behaviour. Although in the future it might be possible to compose several of these models to simulate group interaction, it remains unclear how this might be achieved in practice. I am also unaware of any attempts to apply this approach during an incident investigation.

8.2.4 Hybrid Simulations

This second half of this chapter has identified a number of different simulation techniques that investigators can use to reconstruct the events leading to safety-critical incidents. In identifying these different approaches it has, however, been necessary to introduce what are often arbitrary distinctions. For example, subjunctive simulations can be used during the initial stages of an investigation. Once analysts have reconstructed a likely course of events, this may then be used to produce animated simulations. These prevent viewers from exploring the alternative options that are available to incident investigators. In contrast to subjunctive techniques, users are simply presented with a linear sequence of images that illustrate a particular perspective on an incident. There are further complexities. For instance, declarative reconstructions can be used to model the physical environment in which an incident occurs. However, they provide little information about the way in which particular events contribute to an adverse occurrence. They must be used in conjunction either with text-based time-lines or with abstract simulations of critical events. It can, therefore, be argued that most reconstruction tools exploit hybrid combinations of the techniques that we have reviewed.

Figure 8.16 provides an example of a hybrid approach to incident reconstruction. This design was developed by Gilles Le Galo as part of a proposed simulation tool for incident investigation throughout European air traffic control [423]. As can be seen, the top portion of the display uses data logs to provide real-time reconstructions of the controllers primary information displays. Below this there is an area that reproduces the flight strips that were being used by the controller while they were interacting with this information. This is an interesting use of hybrid reconstruction

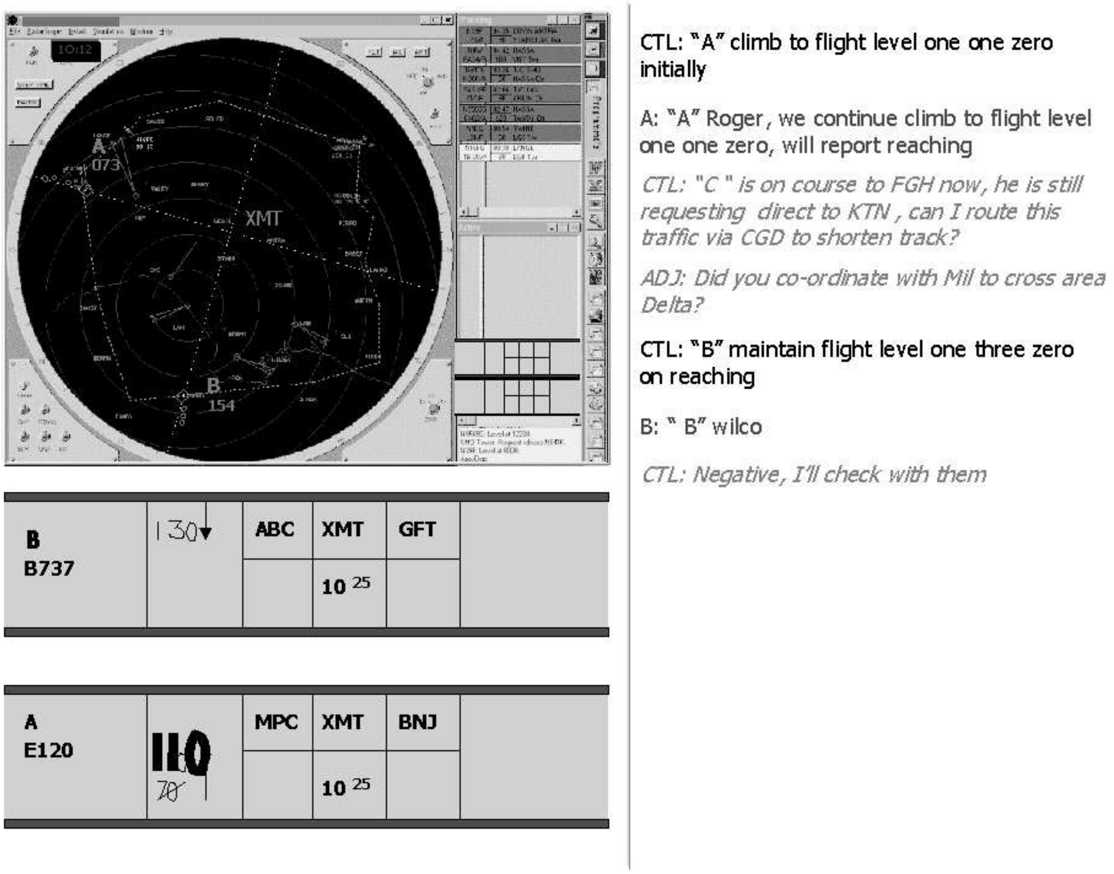

Figure 8.16: EUROCONTROL Proposals for ATM Incident Simulation

techniques because it contains elements of model based simulation; through the regeneration of the primary information displays. It also supports the electronic reconstruction of the controllers' physical environment by presenting scanned images of the relevant flight strips. Finally, the transcript on the right is used to keep track of the radio communications between controllers and the aircrews in their sectors. The vertical, linear presentation of this information is reminiscent of the time-line annotations presented in previous sections.

[212]

There are numerous further examples of this hybrid approach. Various simulation techniques can be supported within a single system, as shown in Figure 8.16. It is, however, more common to use a range of different simulation tools to satisfy different modelling requirements during an incident investigation. This is shown by the range of approaches that are illustrated in Figure 8.17. The image shows a number of different approaches that together contributed to a US National Crash Analysis Center study of ankle injuries in automobile incidents [212]. The image on the top left of Figure 8.17 is derived from a finite element model of joint behaviour. This mathematical technique can be used to account for aspects of joint behavior, muscle tensioning, and injury potential in a high speed impact. The second image on the top of this figure illustrates the process of direct physical measurements that were used to construct a model of the environment in which an injury might occur. These measurements were used to produce the wireframe model on the top-right and the rendered image on the bottom-right of Figure 8.17. Finally, the image on the bottom left shows how a physical model of the lower leg can be used to analyse the causes of injury in an incident. Such direct measurements can be used to validate computer-based simulations of an incident. Mathematical models, rendered visualisations, direct measurements and physical reconstructions all contribute to

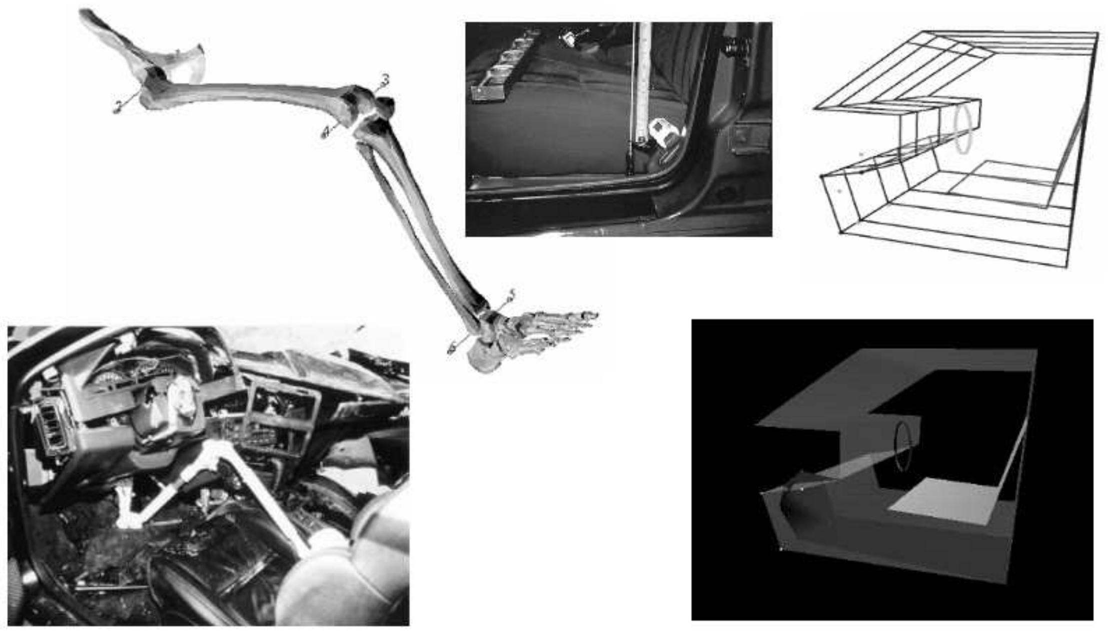

Figure 8.17: US National Crash Analysis Centre's Simulation of Ankle Injury in Automobile Accidents

increase confidence in the investigators' findings. The authors describe this integration of detailed incident investigation techniques and computer modelling as 'unconventional'. It is also, arguably, essential if we are to validate the products of computer-based simulation.

This integration of computer-based modelling and incident analysis has uncovered a number of important results. For example, the National Crash Analysis Centre's study helped to highlight the importance of muscle tension in lower leg injuries. The simulation predicted that greatest injuries occurred when the leg muscles were relaxed, as might be the case if an individual was not anticipating an incident. However, when the leg is tensed it can act like a stiffened beam so that the greatest load must be absorbed by the weakest element. In consequence, severe ankle injuries are likely if the driver is aware of a potential impact. As we have seen these results were both informed and validated by incident investigations. This could not have been done using other forms of reconstruction. "Dummies require adjustable joints and repeated recalibration while cadavers require some form of joint locking device." [212]

This chapter has focussed on computer-based simulations as a means of reconstructing safety-critical incidents. This is justified by the increasing use of these tools in many different application domains. As systems are developed to support the analysis of increasingly complex accidents, the same technology is gradually also being recruited to support the reconstruction of near-miss incidents. It is important to emphasise, however, that many of the techniques embedded within computer-based reconstruction tools are extremely general. It is possible to exploit similar simulation techniques using more conventional, pencil and paper based approaches. For instance, Figure 8.18 illustrates a hybrid approach to incident reconstruction [48]. This is based on materials that were presented by investigators from Australia's Marine Incident Investigation Unit. An engineering model of the piping system for the on-board generators can be printed together with an overview of the layout of the vessel and direct photographic evidence about the state of key components within the system.

It is important to emphasise that paper-based simulations tend to focus on declarative approaches, such as the maps and diagrams shown in Figure 8.18. It is less easy to reproduce the dynamism that characterises both subjunctive and animated reconstructions. There are, however,

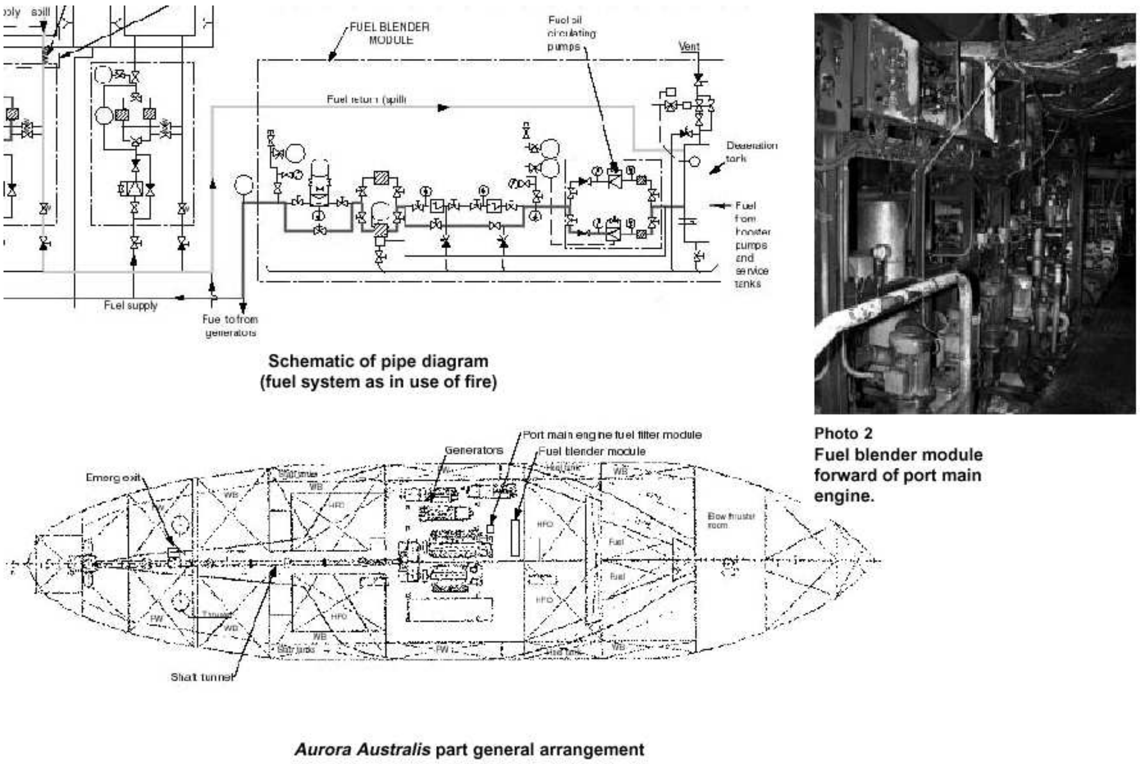

Figure 8.18: Integration of MIIU Plans, Models, Maps and Photographs

some notable exceptions to this general remark. For instance, Figure 8.19 illustrates how the NTSB have used images from desktopVR simulations to illustrate paper-based documents [598]. This image also demonstrates that many previous distinctions between paper-based and electronic reconstructions are becoming extremely blurred. Figure 8.19 shows the electronic version of a paper based incident report. It can be downloaded from the NTSB's web site (www.ntsb.gov). This electronic version of a paper based report was, in turn, generated using an electronic simulation. This translation between electronic and paper-based media is currently not well handled by many investigation agencies. It produces numerous ironies and inconsistencies. Readers who access the NTSB's website can only view the static image that is presented in the paper-based report. This misses an important opportunity to exploit the presentation techniques that are supported by computer-based, simulations.

8.3 Summary

This chapter has introduced the distinction between reconstruction and causal analysis. The former is intended to identify what happened during an incident by simulating the flow of events that led to a near-miss. The latter is intended to explain why those events occurred in the first place. The previous sections, and the next chapter, focus on reconstruction. The intention is to provide a broad overview of the tools and techniques that can be used to piece together evidence for an investigation to develop a coherent account of the course of an incident. This chapter has focuses on computer-based simulations. This is justified by the increasing use of these systems to support incident and accident investigations in many different application domains. In contrast, the next chapter looks more narrowly at the use of textual and graphical notations to model the events leading to failure. These techniques do not require computer support.

It may seem paradoxical to discuss computer-based simulations before paper-based techniques.

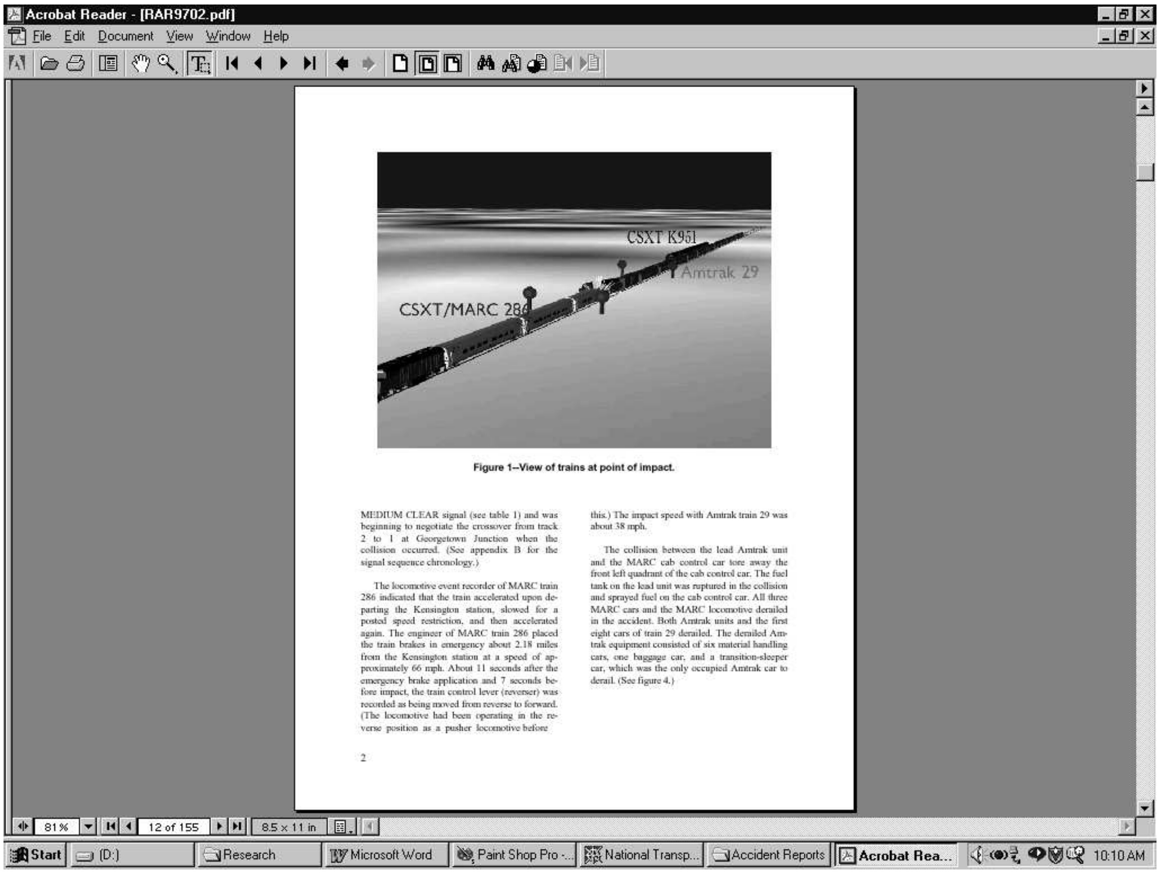

Figure 8.19: NTSB Use of Simulations in Incident Reports

The intention is, however, to first show what is possible using existing software and then to explain why many investigators retain more 'conventional' techniques. Many of the systems that have been introduced in this section are still not widely used. There are many reasons for this and these can be summarised as follows:

- *The financial costs and training overheads are prohibitive.* Many local and regional reporting systems lack the resources to invest in computer-based simulation techniques. There are insufficient funds to acquire and maintain the necessary hardware and software infrastructure. In consequence, many computer-based tools are never applied beyond demonstrator projects. Those that do are most often applied within accident investigations where public and regulatory concern motivate greater investment. These objections are, however, losing their weight. Many simulation techniques, such as QuicktimeVR and the model-based VRML, have been derived from low-cost mass market applications. They can be develop and run on standard PCs and the software is designed to be used by members of the general public;

- *Individual and organisational opposition to innovation.* The reluctance to exploit computer-based simulations can perhaps be explained by the natural conservatism that characterises many investigators. Others have chosen to use more provocative terms. The Rand report has pointed out that having relatively lengthy experience requirements provides strengths and weaknesses for incident investigation [482]. On the one hand, it helps to ensure a detailed first-hand knowledge of an application domain. On the other, it can create a reluctance to exploit novel technologies. This may also apply at an organisational level. The same report criticises the NTSB's lack of any coherent plan to exploit technological innovation in accident and incident investigations. Such comments provoke strong reactions. The Rand report's comments are undoubtedly true for many investigator bodies. It seems paradoxical, however, that some of these criticisms should be applied to an organisation that has also shown a strong desire to innovate. Many of the simulations in this chapter have been inspired by the NTSB's work in this area;

- *Simulations forcer investigators to consider a small subset of incident scenarios.* There are a range of more theoretical objections to the introduction of computer-based simulations. For instance, the visual impact of many computer-based simulations can have an unwarranted effect upon some investigators. This is particularly important given the difficulty of validating the backwards reasoning that drives many simulations. Investigators are often required to estimate the forces or other properties necessary to incur the consequences that are observed in the aftermath of an incident. Unfortunately, the same observed set of outcomes can be derived from many different precursors. This creates particular concerns if simulation tools focus investigators' upon a small number of these potential scenarios.

- *Simulations can encourage encysting and may limit the scope of an investigation.* By 'encysting' we mean that investigators may spend so much of their available resources on developing an elegant simulation that they neglect other aspects of their analysis. This effect is compounded by the fact that some aspects of an incident can be extremely difficult to reconstruct using current tools. For example, it can be hard to reconstruct some meteorological conditions using many existing systems. Similarly, it is hard to envisage ways in which current tools can be used to animate the events leading to managerial and regulatory failures;

- *Reconstructions can raise as many questions as they answer.* Ultimately, simulation tools are only useful if they support the wider objectives of the people and organisations who use them. There is, arguably, a perception that investigators are adequately supported by the pencil and paper-based techniques that are discussed in the next chapter. If Perrow [677] and Sagan [719] are correct then this situation will change as incidents and accidents reflect the increasing complexity of modern technology. It does not, however, follow that computer-based simulations provide an appropriate alternative to these conventional techniques. For instance, current text-based time-lines enable investigators to work at a level of abstraction that is appropriate to the stage of their investigation. Crude sketches, which are produced during an initial enquiry, can be developed and extended as more information becomes available. This 'principle' of proportionate effort is often violated by many simulation tools. Investigators are often forced to accept a range of default parameters because they have no available data about particular aspects of an incident. Some automobile simulators require information about whether lights and wipers were working at the point of impact;

- *Regulatory guidelines can restrict the use of simulation tools.* Previous paragraphs have referred to the 'god's eye' view that is provided by some computer-based systems. These tools enable investigators to integrate data from diverse sources, many of which could not have been accessed by operators during an incident. The resulting animations can have an insidious effect. With the benefit of hindsight and with access to these additional data sources they can used to suggest that operators should have been better prepared for an incident. Reasonable concerns over this mis-use of technology have persuaded some regulators to restrict the use of computer-based simulations [423]. Typically, these only apply to the public dissemination of any resulting animations or models rather than to their use in the reconstruction of an incident.

These objections must be balanced against the benefits of computer-based simulations. These can be summarised as follows:

- *Simulations enable investigators to reconstruct the environment in which an incident occurs.* As we have seen, a range of model-based and photorealistic techniques can be used to reconstruct the layout of a working environment. Investigators can then use interpolation and rendering software to move within those environments. This approach can be used to record the aftermath of an incident, for instance using QuicktimeVR techniques. It can also be used to recreate what the operators might have seen from a number of different locations. The same approaches can also be used to reconstruct particular items that were involved in an incident. A common strength of all these visualisation techniques is that enable users to survey sites that may be too hazardous or expensive for them to visit on subsequent occasions.

- *Simulations enable investigators to replay events in real and virtual time.* The opening sections of this chapter stressed the role of reconstructions in establishing agreement over the course of an incident. Computer-based animations offer considerable flexibility in the way that investigators can play and replay particular events. Software support can be used to alter the real-time of key failures. The sequence in which they occur can also be reviewed and amended. Although there may be initial costs in terms of the time taken to learn how to exploit this authoring software, these are more that off-set by the flexibility of the resulting simulations. In particular, computer-based simulation tools can significantly redice the complexity associated with the maintenance of paper-based documents capture many hundreds of events;
- *Simulations enable investigators to integrate multiple data sources.* Primary and secondary investigations are intended to secure incident data from many diverse sources. Computer-based simulations provide means of integrating this information into hybrid reconstructions. There are many existing problems. For instance, it can be difficult to integrate continuous and discrete data. It can also be difficult to determine the best means of exploiting probabilistic information. Current research continues to explore means of representing operator intervention. In spite of these caveats, computer-based simulations arguably provide the greatest hope of unifying the increasing mass of information that can be obtained in the aftermath of many safety-critical incidents;
- *Simulations enable investigators to explore 'subjunctive' behaviours.* There is an increasing recognition that advanced mathematical techniques can be used to model effects that are difficult or impossible to assess using other approaches. For instance, the effects of muscle tension during traffic accidents can be tested using dummies or cadavers. In the former case, there are considerable recalibration problems associated with the linkages between simulated muscle groups. In the latter case, additional support structures must be used to maintain posture prior to the simulated incident. Alternatively, computer-based models can be derived to provide a low-cost means of simulating the consequences of crashes again and again and again. The results of these studies can be validated using more conventional techniques. However, the costs of crashing full-scale replicas would prohibit the range and scope of tests that can easily be conducted using software environments. An important benefit of these approaches is that support subjunctive simulation. In other words, the low-costs associated with running a trial can encourage investigators to consider a wide range of 'what if' scenarios. This, in turn, encourages the generalisation that was emphasised in the opening sections of this chapter. It is possible to llok beyond the specific events of a near-miss incident to examine potential ways in which such failures might have had more serious consequences.

The use of computer-based reconstruction in incident and accident investigation has had a number of notable successes [9, 212]. It is important to emphasise, however, that many questions remain to be answered about the pragmatic application of these techniques. The development of reconstruction tools to support incident reconstruction lags behind other applications of this technology. The US Aviation Safety Research Act of 1988 provides an insight into the reasons for this lack of investment. This act charged the FAA to undertake "...a research program to develop dynamic simulation models of the air traffic control (ATC) system which will provide analytical technology for predicting airport and ATC safety and capacity problems, and for evaluating planned research projects". The US National Simulation Capability Program was established in response to this congressional mandate. Its goals and objectives reflect the intentions of the Act. Simulations are intended to support training and assist in the development of new systems. The Act arguably neglects the role that simulations can play in understanding the causes of past failures.