

Extending the Application of Formal Methods to Analyse Human Error and System Failure During Accident Investigations

C.W. Johnson and A.J. Telford,
Glasgow Accident Analysis Group,
Department of Computing Science,
University of Glasgow, Glasgow, G12 8QQ.
Tel +44 141 330 6053
Fax +44 141 330 4913
Email fjohnson, alastair@dcsg.gla.ac.uk
<http://www.dcs.gla.ac.uk/~johnson,alastair>

Abstract

Recent disasters at Bhopal, Chernobyl, Habsheim and Kegworth illustrate the point that software is rarely the sole cause behind a major accident. Operator intervention, hardware faults, weather conditions and malicious acts all combine to create the conditions for failure. In the aftermath of these accidents, it seems difficult for software engineers, systems developers, forensic scientists and interface designers to predict all of the ways in which systems can fail. It is, therefore, important that we learn as much as possible from those failures that do occur. Unfortunately, it is often difficult to gain a coherent overview from the mass of detail that is typically contained in many accident reports. This makes it difficult for readers to identify the 'catastrophic event' that produced the necessary conditions for disaster. This paper argues that formal specification techniques can be used to resolve these problems. In particular, Lamport's Temporal Logic of Actions is used to build a unified account of the human errors and system failures that contributed to the Three Mile Island accident. This notation provides high-level abstractions that can be used to strip away the mass of irrelevant details that often obscure important events during disasters. Formal proof techniques can then be applied to the models as a means of identifying the causal relationships that must be broken in order to prevent future failures.

Keywords: Accident analysis; formal methods; human factors; Temporal Logic of Actions.

1 Introduction

Accident analysis is a multi-disciplinary activity [17, 25]. Forensic scientists, metallurgists, meteorologists as well as software engineers and human factors experts all contribute to accident investigations. The reports that are produced by these groups are typically divided into a number of chapters. Each of these sections reflects the perspective of a 'parent' discipline. This makes it difficult for readers to obtain a clear and coherent overview of the events leading to a major accident. Critical failures in one chapter are often missing in other sections. The importance associated with a particular event also varies within many accident reports. This can prevent readers from gaining an accurate impression of the primary and secondary causes leading to human and system failure [9, 10]. Unless regulatory authorities and companies are

convinced about the accuracy of accident reports, then it will be difficult for them to accept the validity of their recommendations.

1.1 The Role Of Formality In Accident Analysis

This paper argues that the application of mathematical specification techniques can be extended from the fields of software and hardware engineering to describe the sequences of events that lead to a major accident. We intend to use formal representations to identify safety requirements after a major accident. This is markedly different from conventional approaches in which fault-trees and other formal notations are used to analyse failures for a system built a central objective in this work is to develop a common framework that can be used to analyse both system and human failures. This approach is appropriate because a number of commercial and regulatory organisations have mandated the use of formal notations in the engineering of safety-critical systems [3].

Human error has played an important role in the course of many major accidents. It is, therefore, important that accident investigation techniques provide means of representing and reasoning about operator intervention. Unfortunately, mathematical specifications have not previously been used to support such analysis. They have, however, been used to support the design of interactive systems. For instance, Harrison and Thimbleby recruited geometric notations to specify high level requirements for interactive systems [4]. Dix has extended this work to investigate menu problems in distributed multi-user systems [5]. Sufriani and He have exploited the Z schema calculus to represent interaction with a text editor [21]. The following pages build on this work to show how formal specification and proof techniques can be extended from the field of software engineering and interface design to support the analysis of human factors and system failures.

Figure 1 presents an overview of the approach that is advocated in this paper. Traditional accident analysis techniques are employed by domain experts, human factors specialists, forensic scientists, software engineers and so on. The aim of this analysis is to identify the causal events that lead to failure. Their findings are typically, expressed in natural language and they form

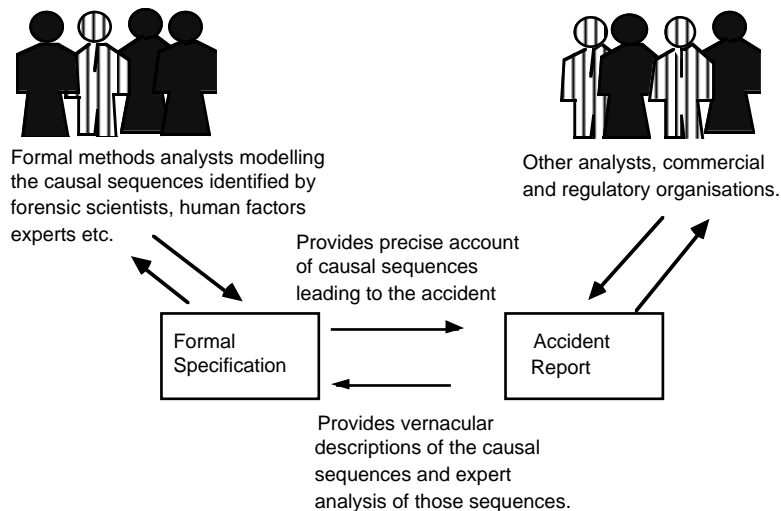


Figure 1: The Role of Formalism in Accident Analysis.

the basis of an initial accident report [1]. This paper argues that mathematical notations can

they be used to model the causal sequences that are identified in these reports. The formalisation process helps to extract these critical sequences from the mass of background detail that may be present during an investigation. Essential information must be represented in the model, contextual detail is omitted. A further benefit is that the resulting models are amenable to formal proof techniques. For instance, an analyst can establish whether or not the causal sequences in a model actually lead to the catastrophic event that occurred during an accident. Finally, the formal model can be translated back into other natural language documents that communicate recommendations to commercial and regulatory organisations.

1.2 Graphical or Textual Notations?

The overview shown in Figure 1 raises a number of practical problems. For instance, there are no well-developed procedures for deriving formal models from the informal accounts in accident reports [9, 24, 23]. Similarly, there has been little work into the utility of formal notation for accident analysis. This is a significant omission because the choice of notation profoundly affects a non-formal analyst's understanding of a formal analysis [12]. For example, Figure 2 shows how the graphical Petri Net notation can be used to model the events leading to the Kegworth accident. This diagram was produced as a result of an earlier collaboration with a multi-disciplinary team of psychologists, linguists and systems engineers. The use of a graphical notation helped the team to discuss the causal sequences that were represented in the model [13]. It also helped to focus our analysis upon particular interpretations of the events leading to failure. For instance, the development of the model forced the team to address the debate about what the First Officer observed on the Airborne Vibration Monitor.

Unfortunately, Figure 2 also illustrates some of the problems that arise with a graphical notation. They quickly become intractable as more and more events are represented. The full diagram for the Kegworth accident takes many pages even though the model is at a relatively high level of abstraction. Such problems can be reduced by structuring mechanisms. For instance, hierarchical Petri nets enable analysts to construct sub-networks, sub-sub-networks and so on. The problem with this is that important information may be hidden within the various layers of the model. A second issue is that the intuitive appearance of graphical notations must often be balanced against correspondingly more complex proof techniques. For example, the vector analysis that is used to establish liveness and safety in the sub-networks of a Petri net cannot easily be explained in terms of intuitive graphical models. This raises a third problem. Analysts may believe that they understand the consequences of introducing a new element into a graphical notation without fully comprehending the additional obligations that are created by their actions. For example, the introduction of an inhibitor arc into a Petri Net can invalidate previous liveness results. Such problems justify the approach described in Figure 4. The 'specification' of the accident is not directly accessed by non-formalists. In contrast, the model is built and interpreted by trained analysts in much the same way that forensic scientists build and interpret continuous models of physical processes such as the Health and Safety Executive's combustion simulator [5]. The U.K. Engineering and Physical Sciences Research Council has recently funded a three-year project that is addressing these issues in more detail (GR/K55040). In anticipation of the results of this work, the remainder of this paper applies textual notation to analyse the events leading to the Three Mile Island accident.

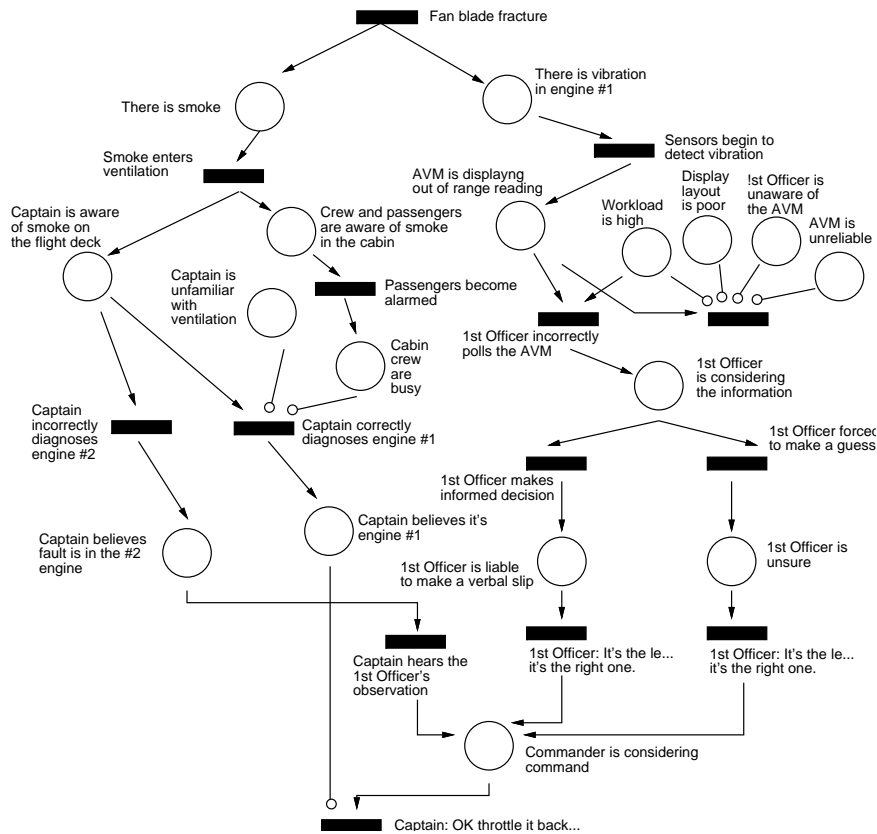


Figure 2: The Events Leading To The Kegworth Disaster.

1.3 Temporal Logic of Actions

The following pages exploit Temporal Logic of Actions (TLA) [16]. This textual notation offers a number of benefits for accident analysis. It provides well developed proof techniques. It also provides means of representing the temporal properties that frequently have a profound impact upon the course of an accident. The underlying model in TLA is formed from an infinite series of states, known as behaviours. For example $\langle\langle s_0; s_1; :: : \rangle\rangle$ denotes a behaviour whose first state is s_0 , its second state is s_1 and so on. From this we derive the definition for \Box (reads 'always') given any formula in the logic F :

$$\langle\langle s_0; s_1; :: : \rangle\rangle \models \Box F \iff \forall n \in \mathbb{N} \text{ at } \langle\langle s_n; s_{n+1}; s_{n+2}; :: : \rangle\rangle \models F \quad (1)$$

This operator can be used to represent invariant conditions that do not change during the course of an accident. The \exists (reads 'eventually') operator can be defined as:

$$\exists F \iff \exists n : F \quad (2)$$

This operator can be used to describe properties that occur at an indeterminate point during interaction. This is a significant advantage because given limited sensing technology and partial sources of evidence it is frequently impossible to develop a precise timeline for all of the events during an accident. The important point here is that analysts can use TLA operators such as \exists , to be precise when they have the necessary evidence. Non-deterministic operators such

as 3, can be used when there is limited information about the temporal ordering of particular failures during a major accident.

The 2 and 3 operators are common features of many different temporal logics [18]. Lamport's Temporal Logic of Actions is distinguished from these notations in that elementary temporal formulas can be used to represent actions. An action A , is a relation between states A . A boolean value can be obtained from applying the action to a pair of states, $[A]$. In other words the action connects the old states, to the new state, by updating elements of that state. Lamport [18] provides additional information on the theoretical foundations of the notation. In contrast, the remainder of this paper focuses upon the application of TLA to support accident analysis.

1.4 The Three Mile Island Accident Overview

The Three Mile Island Accident provides an appropriate example for this paper because it illustrates the pathological mix of system failures and operator error that pose the greatest challenge for software engineers and human factors analysts. Figure 3 presents a schematic overview of the main components in this accident. The problem in this reactor started when the main feedwater system malfunctioned. This was the result of a demineraliser failure. The auxiliary feedwater supply was supposed to have started automatically. It did not because several manual valves in the auxiliary system had been inadvertently left closed after a test of the system. Without the feedwater, the steam generator dried out. This deprived the reactor of its main heat sink and so there was a rise in temperature and pressure inside the primary coolant circuit. The rise in pressure automatically shut-down the reactor, but the radioactive core continued to give out heat. The pressure relief valve opened and became stuck, so that there was a continuous release of radioactive water from the main steam pressure vessel to a drain tank and the containment sumps. Within two minutes the emergency coolant system started automatically and began to increase the coolant level. This originally was of coolant accident

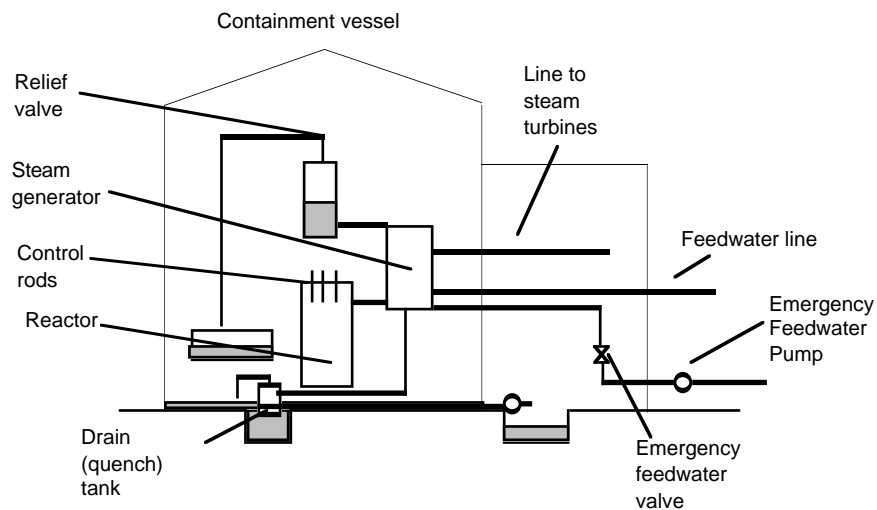


Figure 3: Diagram of the Three Mile Island Pressurized Water Reactor.

(LOCA) was exacerbated by the operators who responded to the pressure warning by manually overriding the two emergency core cooling injection pumps. By fifteen minutes into the accident,

there were almost 40,000 liters of primary coolant on the containment floor as the liquid began to escape from the sumps. Radioactive isotopes of krypton, iodine and xenon escaped to the environment through the building ventilation system.

1.5 Outline Of The Paper

This section introduces the argument that is presented in this paper. Section 2 builds on this and describes a high level framework that can be used to represent and reason about the operation of complex interactive systems. This work has direct parallels with the abstract model that Lamport and others have developed to support the software engineering of concurrent programs. Section 3 then goes on to show how Lamport's model can be used to reason about human error and system failure. Unfortunately, this high level model is too abstract for accident investigations. Section 4, therefore, demonstrates that additional detail can be introduced to represent the specific human factors and system problems that lead to particular accidents. It is then possible to identify 'catastrophic' subsets of these problems which are the primary causes of an accident. Section 5 shows how these can be drawn from the causal sequences that are described in Section 4. It is the identification of these catastrophic failures that traditionally forms the heart of any accident investigation. The important point here is that the use of a formal notation provides a precise means of representing the causal relationships between failures that are often buried within the prose of natural language reports. Section 6 then goes on to demonstrate an additional benefit to four approaches. Formal proof techniques can be applied to TLA descriptions in order to prove that the causal relationships hold given a known set of failures and a normative model of system operation. In other words, we can validate a model of system failure by comparing the causal links in our specification to the known outcome of system failure and operator error. Section 6 summarises the conclusions from this work and suggests areas for future research.

2 A High Level Model For Accident Investigations

There are a number of key aims for any accident investigation. In particular, analysts must identify:

- + the starting point or initial circumstances for any accident
- + the normative or anticipated behaviour of the system
- + the abnormal event that led to system failure and operator error!

This section shows how a model developed to support the software engineering of concurrent programs can also be used to provide a generic framework for these different activities. Subsequent sections will show how detail can gradually be added to this abstract model in order to identify the ways in which critical events combined during a major accident.

2.1 The Elements of Lamport's Model

Many similarities can be drawn between the development of concurrent programs and the generation of accident reports. Both activities are concerned with the analysis of synchronous interaction between complex autonomous processes. In software engineering, designers must account for the behaviour of programs interacting through messaging mechanisms or remote procedure calls. In

the case of accident investigation analysts must examine the ways in which multiple operators interact with a range of hardware and software systems. These similarities are also apparent if one compares the aims of accident investigations, illustrated by the previous list, to the formal component of Lamport's model for the software engineering of concurrent programs [16]. This model describes the operation of a system. Operating conditions in terms of some initial conditions, initial conditions, some changes to the state of the system, $\Sigma \rightarrow \Sigma$, and a fairness condition;

$$\text{Operating conditions} \wedge \text{Initial conditions} \wedge \Sigma \rightarrow \Sigma \wedge F \quad (3)$$

For our purposes an accident report can be thought of as the Operating conditions except that there we are describing a condition of failure rather than a condition of the successful execution of a concurrent program. These reports must represent the initial conditions that hold before an accident occurs, Initial conditions. They must also consider the behaviour of the application during the accident, $\Sigma \rightarrow \Sigma$. Here Σ represents the state of the accident components in terms of a number of variables. Finally, the formal model must ensure that each action had an equal opportunity to occur; otherwise we would not model accurately the progression of events representing the incident. This is captured by the fairness condition, which applies to all the component actions of the next-state relation.

The strength of TLA with regard to accident analysis is that it enables refinements of sub-components of the accident scenario to be developed as more information becomes available to accident investigators. For example, even though the demineralizer of the plant is not mentioned in Bignell and Fortune's account, it would be possible to add a TLA description of it at a later stage, if required. In addition, TLA allows an accident scenario to be described at both the temporal behavioural level and at the functional system level. The actions of the formalism are first-class processes which cannot delay at the level of detail required, the changes that occurred within various components of the disaster. We can, for instance, describe both simple temporal relationships, such as water in the air-circuits leading to some valves switching off and the levels of temperature and pressure in the primary circuit surrounding the reactor core. For this reason of flexibility we have chosen to use TLA over a simple predicate logic.

2.2 State Transition ($\Sigma \rightarrow \Sigma$)

In software engineering, designers must establish that certain output relations will eventually hold given a set of pre-conditions, initial conditions, and some well-defined state transitions, [16]. In accident analysis, investigators must show how an accident eventually occurred given the basic operating conditions of the nuclear plant, Σ , and the set of initial events described in their report, Initial conditions. This parallel between software engineering and accident analysis can be used to structure the formal analysis of our Three Mile Island example. For example, it is possible to introduce a causal relationship, $A \rightarrow B$ (reads 'leads to') that indicates the intended impact of a sequence of instructions in a concurrent program. These are the transitions described by $\Sigma \rightarrow \Sigma$. This same relation can also be used to describe the cause-consequence chains that lead to a major accident. The following formula can be read as it is always the case that if A occurs, then eventually B will happen:

$$A ; B \stackrel{4}{=} \Sigma \rightarrow \Sigma \wedge B \quad (4)$$

It should be noted that the \rightarrow operator formalises a relatively primitive notion of causality. It does not distinguish between necessary and sufficient conditions nor does it describe a struc-

ture for the events in causal sequences. Ladi kind describes a number of more complex logics that can be used to capture detailed notions of causality [15]. Rather than adopt these models, this paper exploits the simple formalisation, to describe the relations between critical failures and an eventual accident. In the following, Operating Conditions represent the conditions existing at the plant, prior to and during the accident. This was defined in (3). Catastrophic Conditions represent the set of failures that occurred on 28th March 1979. These form the traditional basis of accident analysis:

$$\text{Operating Conditions} \wedge \text{Catastrophic Conditions} \wedge \text{Loss Of Coolant Accident} \quad (5)$$

This section has argued that the similarities between the software engineering of concurrent programs and accident analysis can be used to provide a formal structure for accident reports. Lamport's abstract model for the specification of parallel applications has been used to describe the events leading to a major failure as an extremely high level of abstraction. This is important because it helps designers strip away the clutter of low level details that can frequently obscure critical events during an accident. Unfortunately, such a high level of abstraction is inappropriate for many stages of accident analysis. It is therefore important that human factors analysts and system engineers can gradually introduce additional details as their investigations progress. This process is illustrated in the following section.

The Enabled predicate true for an action A and a state s if and only if it is possible to take an A step starting in that state. That is to say that there exists some other state which is related to s by the action A . This predicate is vital to our formalisation of accident analysis in that we can use it to stipulate which events can occur in the model. In particular, we use it to denote the fact that the catastrophic conditions such as certain valves becoming waterlogged at the TMI plant, could take place.

2.3 Initialisation

A key stage in any investigation is the identification of a starting point for an accident. This may seem to be a trivial issue but an analysis of the reports that were produced in the aftermath of Three Mile Island indicates that this is not the case. For example, our account of the accident in Section 4 followed [7] in stating that the original cause was a demineraliser failure. Other reports blame the accident upon the failure of the air supply to an air operated valve [2]. It is clearly important to reduce such ambiguity if software engineers, hardware developers and interface designers are to coordinate the development of future systems. Formal methods can be used to provide an unambiguous representation of the initial circumstances that surround an accident. For example, the elements of TLA can be used to provide additional details for our formalisation of initial conditions. The following clause reflects the second interpretation mentioned above. It formalises the condition for air supply failure rather than a demineraliser fault. The state of the demineraliser is omitted because in this analysis it is not considered to be a primary cause of the failure. Of course, this information might be included if analysts revised their interpretation of the events leading to the accident.

$$\begin{aligned} \text{Init} = & (\text{relief_valve_closed} \wedge (\text{primary_ow} = \text{normal}) \\ & \wedge (\text{steam_turbine_on}) \wedge (\text{e_pumps_status} = \text{online}) \\ & \wedge (\text{e_pumps} = \text{off}) \wedge (\text{air_circuit_status} = \text{clear}) \\ & \wedge (\text{secondary_ow} = \text{normal}) \wedge (\text{ac_valves} = \text{open}) \end{aligned}$$


```

^ (controlrods= up) ^ (acpumps= on)
^ (waterunofi= false) ^ (watertemp= normal)
^ (waterpress= normal) ^ (fedwaterline open)
^ (reliefvalvestatus working) ^ (hppenable= online)
^ (letdownvalve closed) ^ (highpressurepumps= ofi)

```

(6)

The variables in this clause will be the value of, in \mathbb{N} , at the start of the accident. It should be stressed that this predicate is in itself an abstraction of the mass of more detailed information that may be presented to courts of enquiry or in subsequent accounts [17]. For example, primary flow and secondary flow indicate the flow of the water in the primary and secondary circuit steam turbines as an abstract representation of the state of the steam turbine. The epumps and dpumps status variables indicate whether the emergency feedwater pumps were on and whether they were enabled or not. The explicit representation of initial conditions is an important stage in the development of accident reports just as it is for the engineering of concurrent programs. In both cases, it helps to frame the subsequent design and analysis by providing starting point for further consideration.

2.4 Fairness Conditions (F)

The second element of Lamport's model is the fairness condition. In software engineering terms this is used to ensure that events which must happen, eventually do happen:

$$F = \bigwedge_i (23 \text{ hAi}_f) _ (23 : \text{Enable} \text{ hAi}_f) \tag{7}$$

At first sight, it might seem strange to consider fairness conditions during the analysis of an accident. However, the previous disjunction precisely characterises the aim of accident reports. These documents are intended to describe how an accident eventually occurred so that it can never happen again. Indeed we have strong fairness conditions covering both the ordinary actions of the plant, which we denote C , and the catastrophic actions, denoted as Catastrophic Conditions. This ensures that both the catastrophic actions and the operating actions may occur with equal priority in the model of the event. When we further specify that,

Enable Catastrophic Conditions

we are saying that the latter part of the fairness disjunction is true and, as a consequence,

$23 \text{ hCatastrophic Conditions}$

This says that the catastrophic actions must occur which is naturally consistent with the disaster scenario that we are formalising.

3 Causal Sequences (;)

One of the ways in which formal notations can add detail to high level accounts of accident reports is by representing the causal chains that lead to a major accident. These chains are built from the operating conditions, operating conditions and the fact that the catastrophic events, Catastrophic Conditions, that led to the loss of coolant accident, LossOfCoolant Accident, were enabled. An important technical point is that subsequent sections will define the Enable

predication of their interest of clarity. It is important to have an explicit representation of these sequences because it can be difficult for readers to extract causal relationships from the natural language accounts of a major accident. For example, one report includes the following lines:

:::they [the maintenance crew] allowed some water to enter an air circuit that opened and closed some valves ::: The affected valves shut off ::: ([2], p. 12)

This quotation illustrates the ordering and selection of materials used to imply a causal relationship between water entering the air circuit and the valves being shut off. The following clause makes this relationship explicit: Water in Air Circuit denotes water being introduced into the air circuit and Valve Switch Off stands for the valves switching off in the circuit. This illustrates causal chains reformed from the catastrophic events Catastrophic Conditions such as the introduction of water into the air circuit and from the operating conditions, Operating Conditions that determine the system's response to the failure:

Water in Air Circuit ; Valve Switch Off (8)

Having constructed such a causal sequence of software and hardware engineering can trace the ways that subsequent events stemmed from these initial failures. For instance, the same report states that:

The affected valves shut off and the pumps in the same circuit closed down in quick succession.

The same; relation can be used to represent the consequences of the valve failure which in turn was caused by the water entering the air circuit. In the following clause Pumps Closing Down is the action of the pumps closing down:

Valve Switch Off ; Pumps Closing Down (9)

The important point here is that the; relation forces analysts to consider the ways in which a wider range of failures contribute to a major accident. The initial events might simply relate to the hardware or systems engineering but the knock-on effects of those failures are quickly propagated throughout the application. In our example, the failure of the valves led to the pumps closing down. This led to the feedwater line being closed. The consequences of this was that the automated control systems intervened. They shut down the steam turbines and started up the emergency turbines. As software is increasingly being introduced into such safety-critical applications it is vital that designers understand how such knock-on failures affect the overall requirements for the systems. Unfortunately, it is not easy to extract these requirements from natural language accounts:

One of the pumping circuits affected [by the valve switching off] was the feedwater line to the steam generator in the secondary circuit ::: so the safety systems shut down the steam turbine and the electrical power generator it drove. ::: When the pumps supplying water to the steam generator stopped, three emergency pumps for the feedwater started.

As before, formal notations can be used to focus upon the causal relationships that are embedded within such informal observations. Steam Feedwater Down represents the feedwater line to the steam generator being closed. Steam Turbines Shut Down represents the shutdown of

the steam turbine and Emergency Pumps_Start is the action of the emergency pumps starting up:

Pumps_ClosingDown ; Steam_Feedwater_Down (10)

Steam_Feedwater_Down ; (Steam_Turbine_Shutdown ^ Emergency_Pumps_Start) (11)

Thus far we have shown how a hardware failure led to the interaction of control systems in order to start up emergency pumping equipment. These problems were exacerbated by operator interaction or lack of it:

:::an operator::: noticed that the pumps were running but he did not notice two particular warning lights on a control panel. They signalled that valves were closed on each of the two emergency feedwater lines and so were preventing water from reaching the steam generator, which soon boiled dry::: Not until almost eight minutes later did the operator notice the closed valves and open them.

We use Emergency_Valves_Closed to represent the fact that the emergency feedwater valves were closed and Steam_Generator_Dry stands for the steam generator boiling dry.

(Steam_Feedwater_Down ^ Emergency_Pumps_Start ^ Emergency_Valves_Closed) ;
Steam_Generator_Dry (12)

The absence of operator interaction can be represented by the following clause:

(Steam_Feedwater_Down ^ Emergency_Pumps_Start ^
Emergency_Valves_Closed ^ :OperatorOpensValves ;
Steam_Generator_Dry (13)

Such formulae require further refinement if they are to guide the future activities of interface designers, systems engineers and programmers. For instance, there must be some explanation as to why the operators did not intervene to open the valves, :OperatorOpensValves. The subsequent sections address this issue in greater detail. In contrast, the remainder of this section continues to build the causal chain of events leading to the Three Mile Island accident.

Deprived of an outlet for the heat still entering from the reactor core, the water in the primary circuit increased in temperature and pressure. A relief valve on top of the pressurizer opened and steam and water began to flow out to a drain tank::: To reduce the production of heat in the reactor core, the control rods automatically descended into it:::

In the following, Primary_Circuit_Emergency stands for the emergency conditions which resulted in the primary circuit. In particular, there was a rise in the temperature and pressure in the circuit. Control_Rods_Drop denotes the fact that the control rods dropped into the reactor core and Relief_Valve_Open indicates the opening of the relief valve on the primary water circuit. Relief_Valve_Open represents the relief valve being open and Water_Run_Off denotes water run-off into the drain tank from the primary circuit:

Steam_Generator_Dry ; Primary_Circuit_Emergency (14)

Primary_Circuit_Emergency ; Control_Rods_Drop ^ Relief_Valve_Open (15)

Relief_Valve_Open ; Water_Run_Off (16)

Formula (15) deserves particular attention because it illustrates a causal chain of failures that can have multiple consequences. An increase in temperature and pressure led to the introduction of the control rods and to the opening of relief valves. Such combined actions are often necessary in order to restore normal states:

The combined effects of lowering the control rods and opening the relief valve brought the temperature in the vessel down to normal::

The use of a formal notation would provide few benefits if it could not trace the ways in which automated systems and safety applications successfully intervened:

ControlRods.Drop ^ WaterRun.Ofi ; : PrimaryCircuitEmergency (17)

This illustrates some of the complexities that can arise during accident investigations. It also emphasizes the importance of tracing the full extent of chains. The fact that the temperature and pressure problems were being reduced; PrimaryCircuitEmergency did not mean that the accident was over. One of the relief valves was faulty and did not close as the pressure and temperature returned to normal:

Through this valve more than one third of the content of the primary circuit escaped. :: two high pressure pumps started automatically, triggered by a lowering of pressure in the primary circuit.

The following clause, therefore, states that if there is no pressure emergency in the primary circuit PrimaryCircuitEmergency then water WaterRun.Ofi, run-off will lead to a loss of pressure PrimaryPressureLoss. The second clause represents the observation that the high pressure pumps start HighPressurePumps.Start as a consequence of the loss of pressure if the pumps are enabled HighPressurePumps.Enabled:

: PrimaryCircuitEmergency
(WaterRun.Ofi ; PrimaryPressureLoss (18)

HighPressurePumps.Enabled)
(PrimaryPressureLoss; HighPressurePumps.Start (19)

This final sequence provides the missing link between the initial ingress of water into the air circuit and the eventual loss of coolant in a LOCA:

The let-down valve was opened to release what was believed to be an excess of water. The original fall in pressure and the failure of the injection to bring down the temperature should have alerted the operators that a Loss of Coolant Accident was underway but it did not.

The following clause states that a loss of pressure in the primary circuit together with opening the let-down valve, LetDown.Open, led to the Loss of Coolant Accident:

PrimaryPressureLoss ^ LetDown.Open ; LossOfCoolantAccident (20)

It is important to re-iterate the main point of this section which is that the complexity of natural language descriptions makes it difficult to reconstruct a causal sequence that leads to a major failure. The reader is left to interpret the relationships that are built up from natural language statements such as 'and then', 'at which time', and so... 'in quick succession' and 'they signalled that'. The use of the; relation avoids such ambiguity. It also provides a concise means of representing the human errors and system failures that led to the loss of coolant accident.

4 Catastrophic Events (Catastrophic Conditions)

The previous section has constructed a sequence of events that characterise the course of a major accident. These sequences are built from a number of failures and from the responses of the system and its operators to those failures. Operating conditions having constructed these causal sequences it is possible for analysts to extract those critical failures which were essential for the accident to progress to catastrophic conditions. For the Three Mile Island example, these catastrophic events can be represented as follows:

$$\text{CatastrophicConditions} = \text{WaterinAirCircuit} \wedge \text{ReliefValveFailure} \wedge \text{EmergencyValvesClosed} \wedge \text{LetDown_Open} \wedge \text{EmergencyPumps_Off} \quad (21)$$

The accident was caused by the introduction of water into the air circuit and the failure of the relief valve and by the operators failing to close the relief valve and the opening of the letdown valve in the primary circuit and the emergency pumps being disabled. In the above, EmergencyPumps_Off refers to the emergency pumps being disabled and switched off:

$$\text{EmergencyPumps_Off} = \text{HighPressurePumps_Enabled} \wedge \text{HighPressurePumps_Start} \quad (22)$$

These high level abstractions enable analysts to focus upon critical events. Hardware and software engineers are not forced to consider the mass of circumstances and detail that must be presented to accident inquiries.

4.1 The Causes of Catastrophic Events

Catastrophic conditions is the conjunction of the actions that contributed to the accident taking place. The accident occurred because these actions were enabled in some way. For example, (8) stated that the ingress of water into an air circuit caused the valves to switch off. This in turn led to the pumps being switched off. Unfortunately, cause (8) does not state the reasons why water was allowed to enter into the air circuit. This is a significant omission. Reason [20], Perry [19] and Leveson [17] all argue that analysts must trace the organisational factors that create the 'conditions for disaster'. Such problems can be avoided by extending the causal chain back to the 'primitive conditions for an accident'.

$$\text{AirCircuitMaintenanceError}; \text{WaterinAirCircuit} \quad (23)$$

This raises the problem of when to stop the formalisation process. Additional causes can be introduced to represent the reasons for the maintenance failure; the operator might then be used to represent the causes for the reasons that lie behind the maintenance failure and so on. In practice, the decision to end the formalisation process depends upon the skill and experience of the analyst. The important point here is that the construction of the causal sequence forces analysts to explicitly consider the catastrophic event that frames an accident.

It is also important to note that the elements of Catastrophic Conditions make no distinction as to the source of the 'failure'. In keeping with the title of this paper, we are equally concerned to model both human 'error' and system failure. However, it is entirely possible to represent different categories within Catastrophic Conditions.

$$\text{OperatorError} = \text{WaterinAirCircuit} \wedge \text{EmergencyValvesClosed} \wedge \text{LetDown_Open} \wedge \text{EmergencyPumps_Off} \quad (24)$$

This categorisation might be further refined to identify those events that are the side-effects of operator tasks. Emergency Pumps Off, as well as those that are the result of overt failures, such as Water in Air Circuit or mistakes based upon a misreading of available evidence such as Let Down Open. Such an analysis would be similar to Reason's taxonomy of error [20]. It is also possible to extract those critical events that were due to systems failures:

$$\text{System failure} \stackrel{4}{=} \text{Relief Valve Failure} \quad (25)$$

The previous clauses illustrate the point that TLA does not provide a panacea for accident analysis. It is entirely possible to disagree with our classification of human 'error' and system 'failure'. For instance, the fact that operators could incorrectly open the let down valve might be viewed as a system failure rather than an operator mistake. The notation does not replace the value judgements and skill-based decisions that must be employed to identify critical events from the mass of background detail. Similarly, it does not provide any automatic means of distinguishing the elements of Operator error from System failure.

It may also be noted that there is not universal agreement among those reporting on the accident about its causes. For example, Fremling has cited the possibility of a demineralizer error as the cause of the Three Mile Island accident [7]. This problem occurs within accident reports where different chapters reflect the interests of separate analysts. We investigate inconsistencies between reports in the forthcoming [21]. We are here, however, attempting to formalise and establish consistency within a single report only.

4.2 Refining Catastrophic Events

Clauses such as (21) and (22) are more detailed than the generic model for concurrent programs and accident analysis introduced in (3). Unfortunately, they are still high level, that human factors analysts and programmers might use to guide the development of future applications. This problem can be addressed by representing the changes that critical events cause to process parameters. For example, the introduction of liquid to the air circuit caused ducts to become waterlogged. The value of particular variables, such as air circuit status, can be directly related to the sensor readings that must be polled by control programs:

$$\text{Water in Air Circuit} \stackrel{4}{=} (\text{air circuit status} = \text{clear}) \wedge (\text{air circuit status} = \text{water filled})$$

It is important to emphasise that critical events will typically have no impact upon most of the variables in a system. This can be represented by introducing Unchanged (Var=S), to indicate that all variables in the system not in the set S are unchanged by a specified action:

$$\text{Water in Air Circuit} \stackrel{4}{=} (\text{air circuit status} = \text{clear}) \wedge (\text{air circuit status} = \text{water filled}) \\ \wedge \text{Unchanged (Var=fair circuit status)}$$

The failure of the relief valve was another catastrophic event that contributed to the accident. This event had no impact upon the rest of the system apart from causing the valve to fail. Again, this illustrates the abstraction of formula notation can be used to scope or frame the impact of critical failures:

$$\text{Relief Valve Failure} \stackrel{4}{=} \text{relief valve status} = \text{fault} \\ \wedge \text{Unchanged (Var=relief valve status)}$$

The previous clause represents the impact of a hardware failure. It is also possible to represent the impact of operator error in a similar manner. For example, the disabling of the high pressure coolant pumps played a major part in the events leading to the accident. The pumps were switched from a state in which they were enabled but off-line to one in which they had been turned off:

$$\text{EmergencyPumps_Off} = (\text{highpressurepumps}^0 = \text{off})$$

The fourth of the critical events was the closure of the emergency feedwater valves:

$$\text{EmergencyValves_Closed} = (\text{fedwaterline}^0 = \text{closed}) \wedge \text{Unchanged}(\text{Var}=\text{fedwaterline})$$

The final critical factor in the lead up to the accident was the opening of the letdown valve in the primary circuit. This had the effect of cutting off the primary coolant flow to the reactor:

$$\text{LetDown_Open} = ((\text{waterpressure}^0 = \text{low}) \wedge (\text{letdownvalve}^0 = \text{open})) \wedge \text{Unchanged}(\text{Var}=\text{primaryflow}; \text{letdownvalve})$$

Similar definitions can be introduced for all of the events that have been introduced in our formal model of the accident. This is omitted for the sake of brevity and the interested reader is directed to [14]. In contrast, the following section argues that the formal proof techniques can be used to determine whether causal sequences can be reconstructed from the failures. Catastrophic Condition and the operating procedures in Operating Conditions

5 Proofs and properties

The formal modelling of a major accident provides a number of benefits for human factors and software engineers. It avoids the ambiguity that can be found in natural language descriptions of causal relationships. It provides a precise means of framing an accident. TLA clauses can be used to describe the initial condition for failure, initial conditions. A further advantage is that formal proof techniques can be used to determine whether the causal sequences described by ; can actually be produced from the known operating conditions of the system. Operating Conditions and the catastrophic system failures and human errors in Catastrophic Conditions. For example, accident investigators might want to prove that if the operating procedures at Three Mile Island were followed then the introduction of water into the air circuit would lead to the valves being switched off. If this were not the case then either our causal analysis is incorrect, something else triggered the valves. Or our understanding of the operating conditions is incorrect:

$$\text{OperatingConditions} \wedge \text{EnableWaterinAirCircuit} \wedge \text{ClosedValves} \quad (26)$$

Where ClosedValves is defined as follows:

$$\text{ClosedValves} = \text{acvalves}^0 = \text{closed}$$

As in previous sections of this paper, we can again recruit software engineering techniques to support accident investigation. In this case we can use Lamport's (SF1) proof rule [16] page

22] to help establish the previous relationship:

$$\frac{P \wedge M \wedge (P \supset Q) \wedge (M \supset A)}{2 P \wedge 2 M \wedge 2 F \wedge 3 \text{Enable} \wedge A}$$

$$2 M \wedge SF(A) \wedge 2 F \wedge (P ; Q)$$

The following substitutions can be made between our model of the Three Mile Island accident and the tokens in Lamport's rule:

$$P \wedge \text{EnableWaterinAirCircuit} \wedge M \wedge N \wedge Q \wedge \text{ClosedValves}$$

$$A \wedge C \qquad v \wedge f \qquad 2 F \wedge SF_f(N)$$

The first premise of SF1 cannot be written as follows:

$$\text{EnableWaterinAirCircuit} \wedge N \wedge f \wedge (\text{EnableWaterinAirCircuit} \wedge \text{ClosedValves}) \quad (27)$$

This is satisfied trivially since, in this case, if EnableWaterinAirCircuit is true then EnableWaterinAirCircuit must automatically be true. The second premise can be re-written as follows:

$$\text{EnableWaterinAirCircuit} \wedge N \wedge C \wedge f \wedge \text{ClosedValves} \quad (28)$$

This is true since we are assuming that the initial conditions, InitialConditions, hold. Consequently it must follow that an air circuit is status water filled and subsequently from WCV that ClosedValves is true. The final premise can be written as:

$$2 \text{EnableWaterinAirCircuit} \wedge 2 N \wedge f \wedge SF_f(\text{EnableWaterinAirCircuit}) \wedge 3 \text{Enable} \wedge C_f \quad (29)$$

The above result from the fact that WCV is part of the disjunction which comprises InitialConditions. Consequently since we assume here that EnableWaterinAirCircuit is always true then WCV must be enabled at some point. By applying SF1 we then obtain the following:

$$\text{InitialConditions} \wedge (2 N \wedge f \wedge SF_f(C) \wedge SF_f(\text{CatastrophicConditions}))$$

$$(\text{EnableWaterinAirCircuit} \wedge \text{ClosedValves}) \quad (30)$$

Finally, we observe that (30) is isomorphic to

$$\text{OperatingConditions} \wedge (\text{EnableWaterinAirCircuit} \wedge \text{ClosedValves}) \quad (31)$$

Such proof techniques suggest further ways in which our application of Lamport's model might benefit from software engineering techniques. For example, the transitivity of the ; relation suggests that certain sub-claims in our causal sequence might be omitted. In other words, if A ; B and B ; C then it may be possible to omit B altogether. An example of this may be seen in the above formalisation where the action of closeValves and the variable c_valves could be omitted. The action of the air circuit pumps switching off would then follow directly from the water being introduced into the air circuit. Such formal analysis might then be used to remove unnecessary parts from an accident report. This is similar to the ways in which pieces of redundant or irrelevant code may be identified and removed through mathematical analysis. However, this proposal also illustrates some of the dangers that arise when transferring software engineering techniques directly to accident analysis. The purpose of writing an accident report is not to produce a minimal account but rather to produce a document that effectively communicates the weaknesses of previous systems. Removing such transitive relations may reduce the volume of an accident report but it can also hide links that help readers to understand critical sequences of human error and system failure.

6 Conclusions and Future Work

This paper has argued that Lamport's model for the software engineering of concurrent systems can be used to support the generation of accident reports. This provides a number of benefits. The elements of Lamport's model can be used to identify generic requirements for accident investigations. In particular, multi-disciplinary teams must agree upon the starting point for an accident. Initial conditions, the anticipated behaviour of a system, operating conditions and the abnormal events that lead to failure are catastrophic conditions. The elements of the model can also be applied to structure the detailed analysis of the events leading to a major accident. The Temporal Logic of Actions notation can be used to construct the causal sequences of human error and system failure that lead to disaster. The ; relation can be used to explicitly present the informal observations that are often buried within informal accounts. Finally, proof techniques can be applied to the elements of the model to determine whether the ; relation actually does explain the circumstances leading to an accident.

The work reported in this paper is the initial product of an Engineering and Physical Sciences Research Council project into the formal analysis of human factors and system failures. Much remains to be done. In particular, a powerful application of our approach is that it can be used to detect inconsistencies not only within but also between accident reports. For example, our description of the Three Mile Island accident has been drawn from a number of sources, including [7] and documents issued by the US President's Commission and the US Nuclear Regulatory Commission. These, in turn, were the distillate of evidence presented to the courts of enquiry. Our formalisation is therefore the result of a series of interpretations and simplifications. However, the techniques that are presented in this paper provide a precise and concise account of the events leading to a major accident. It is hypothesised that such formalisation might help to reduce the potential ambiguities that frequently arise in these different accounts.

A further benefit of formalisation that it also provides engineers and analysts with proof techniques that can be used to verify cross-viewpoint consistency. This again illustrates the wider applications of software engineering techniques for accident analysis. The need for such consistency has been recognised by a growing number of researchers in the field of software engineering [64]. This work is an important starting point for our future research. It is vital that those who are affected by the aftermath of an accident are left with a consistent view upon the implications of a report. This will mean, for instance, that consistency has to be shown between nuclear and software engineers. It is also necessary to maintain consistency between human-computer interaction specialists and those who are interested in the computational aspects of system development.

Acknowledgement

Thanks are due to the referees whose comments helped to shape and focus the final version of this paper. Thanks also go to the members of the Glasgow Accident Analysis Group, to the Glasgow Interactive Systems Group and to the Formal Methods and Theory Group in Glasgow University. This work is supported by the UK Engineering and Physical Sciences Research Council grants GR/J07686, GR/K69148 and GR/K55040.

References

- [1] Air Accidents Investigation Branch, Department of Transport. Report On The Incident Involving BA C 1-11G-AYWB and Boeing 737 EI-BTZ On 12 April 1988 At Gatwick Airport, number 2/89 London, United Kingdom, 1989 Her Majesty's Stationery Office.
- [2] V. Bignell and J. Fortune. Understanding System Failures. Manchester University Press, 1984.
- [3] J. Bowen and V. Stavridou. Safety-critical systems, formal methods and standards. Software Engineering Journal, pages 189-209, July 1993.
- [4] H. Bowman, J. Derrick, P. Linington and M. Steen. Cross-viewpoint consistency in Open Distributed Processing. Software Engineering Journal, pages 44-57, January 1996.
- [5] A. J. Dix. Formal Methods For Interactive Systems. Academic Press, London, United Kingdom, 1991.
- [6] A. Finkelstein and I. Sommerville. The viewpoints FAQ. Software Engineering Journal: Special Issue on Viewpoints for Software Engineering, pages 2-4, January 1996.
- [7] J. H. Fremlin. Power Production: What Are The Risks? Adam Hilger Ltd. Boston, United States of America, 1985.
- [8] M. D. Harrison and H. W. Thimbley, editors. Formal Methods In Human Computer Interaction. Cambridge University Press, Cambridge, United Kingdom, 1989.
- [9] C. W. Johnson. The formal analysis of human-computer interaction during accidents. In G. Cockton, S. W. Draper and G. R. Swire, editors. People And Computers IX. Cambridge University Press, Cambridge, United Kingdom, 1994.
- [10] C. W. Johnson. The application of Petri Net store representation and reason about human factors problems during accident analysis. In P. Palanque and R. Bastide, editors. The Second European Workshop On The Design, Specification And Verification of Interactive Systems, pages 345-357. Springer-Verlag, Berlin, Germany, 1995.
- [11] C. W. Johnson. Representing and reasoning about the impact of environmental layout upon human computer interaction. Ergonomics, 39 (3) : 512-531, 1996.
- [12] C. W. Johnson. Techniques for the evaluation of design notation. In M. D. Harrison and J. Vanderdonk, editors. Design, Specification and Verification of Interactive Systems '96. Berlin, Germany, 1996. Springer-Verlag.
- [13] C. W. Johnson, J. C. McCarthy, and P. C. Wright. Using a formal language to support natural language accident reports. Ergonomics, 38 (6) : 1265-1283, 1995.
- [14] C. W. Johnson and A. Telford. The formal analysis of human factors errors and system failures. Technical Report TR-1996-6, Department of Computer Science, University of Glasgow, Glasgow, Scotland, 1996.
- [15] P. Ladkin. Reasons and causes. Technical report Technische Fakultät Universität Bielefeld, Germany, 1996. Available at <http://www.tu-fak.uni-bielefeld.de/>.

- [16] L. Lamport. The temporal logic of actions. Research Report 79, DEC Systems Research Center, 130 Lytton Avenue, Palo Alto, California 94301 USA, December 1991.
- [17] N. G. Leveson. *Safeware: System Safety and Computers*. Addison Wesley, Reading, United States of America, 1995.
- [18] Z. Manna and A. Pnueli. Verification of concurrent programs: The temporal framework. In R. S. Boyer and J. Strother Moore, editors, *The Correctness Problem in Computer Science*, pages 215-273. Academic Press, London, United Kingdom, 1981.
- [19] C. Perrow. *Normal Accidents: Living with High-Risk Technologies*. Basic Books, New York, United States of America, 1984.
- [20] J. Reason. *Human Error*. Cambridge University Press, Cambridge, United Kingdom, 1990.
- [21] B. Suffriani and J. He. Specification, refinement and analysis of interactive processes. In M. D. Harrison and H. W. Thimbley, editors, *Formal methods in Human Computer Interaction*, pages 153-200. Cambridge University Press, Cambridge, United Kingdom, 1990.
- [22] Alastair Telford and C. W. Johnson. *Viewpoints for accident analysis*. Technical report, Department of Computer Science, University of Glasgow, Glasgow, Scotland, 1996. Submitted to *Formal Aspects of Computer Science*.
- [23] M. Thomas. A proof of incorrectness. *ICP: The editing problem in the Therac-25 High Integrity Systems Journal*, (1) : 49-62, 1993.
- [24] M. Thomas. The story of the Therac-25 in LOTOS. *High Integrity Systems Journal*, (1) : 3-15, 1993.
- [25] D. Woods, L. Johansen, R. Cook, and N. Sarter. Behind human error. *Cognitive systems, computers and hindsight*. Technical Report 94-01, SERIAC State of the Art Report, 1994.