The Application of Causal Analysis Techniques for Computer-Related Mishaps

Chris Johnson,

Dept. of Computing Science, University of Glasgow, Glasgow, G12 9QQ.
Tel.: +44 141 330 6053, Fax: +44 141 330 4913
johnson@dcs.gla.ac.uk

Abstract. Causal analysis techniques support the investigation of incidents and accidents. These include elicitation methods, such as Barrier Analysis, and event-based techniques, for example accident fault trees. Other approaches rely on flow charts, including those within the PRISMA approach and accident models, including the control theory model in STAMP. A further class of causal analysis techniques relies upon models of argumentation, such as the counterfactual approach in WBA. This paper reviews the support that different causal analysis techniques provide for the investigation of adverse events and near misses involving Electrical, Electronic or Programmable, Electronic Systems (E/E/PES). The events leading to an explosion and fires at a fluidized catalytic cracking unit are used to illustrate the application of these different techniques. This is then used to assess the degree of support that different techniques provide for the identification of latent failures at different stages in the software and systems lifecycle.

Introduction

The following pages introduce techniques that investigators can use to identify the root causes of incidents involving Electrical, Electronic or Programmable, Electronic Systems (E/E/PES). We refer to E/E/PES rather than ‘software’ or ‘hardware’ because this is the term adopted by the UK Health and Safety Executive when referring to the broad class of programmable systems that are exploited by the process industries. Causal analysis is a process by which investigators can identify the reasons *why* a mishap occurs. In contrast, mishap reconstruction identifies *what* happened during an accident or incident.

An E/E/PES case study will be used to illustrate the causal analysis techniques in this paper. This incident has been chosen through consultation with the UK Health and Safety Executive (HSE) and industry representatives because it typifies the adverse events that currently threaten many safety-critical industries. The following pages describe an incident involving a fluidised catalytic cracking unit, part of a UK refinery complex. The plant receives crude oil, which is then separated by fractional distillation into intermediate products, including light and heavy diesel, naptha, kerosese and other heavier components. These heavier elements are eventually fed into the fluidised catalytic cracking unit. This is a continuous process to convert ‘long’ chain hydrocarbons into smaller hydrocarbon products used in fuels. The immediate events leading to the incident started when lightning started a fire in part of the crude distillation unit within the plant. This led to a number of knock-on effects, including power disruption, which affected elements in the fluidised catalytic cracking unit. Initially, hydrocarbon flow was lost to the deethaniser, illustrated in Figure 1. This caused the liquid in the vessel to empty into the next stage debutanizer. The control system was programmed to prevent total liquid loss in these stages and so valve A was closed. This starved the debutanizer of feed. The programmable system again intervened to close valve B. The liquid trapped in the debutanizer was still being heated even though both valves now isolated it. Pressure rose and the vessel vented to a flare. Shortly afterwards, the liquid level in the deethaniser was restored, the control system opened valve A and the debutanizer received further flow. Valve B should have opened at this time to allow fluid from the pressurised debutanizer into the naptha splitter. Operators in the control room received misleading signals that valve B had been successfully reopened by their control system even though this had not occurred. As a result the debutanizer filled with liquid while the naptha splitter was emptied.

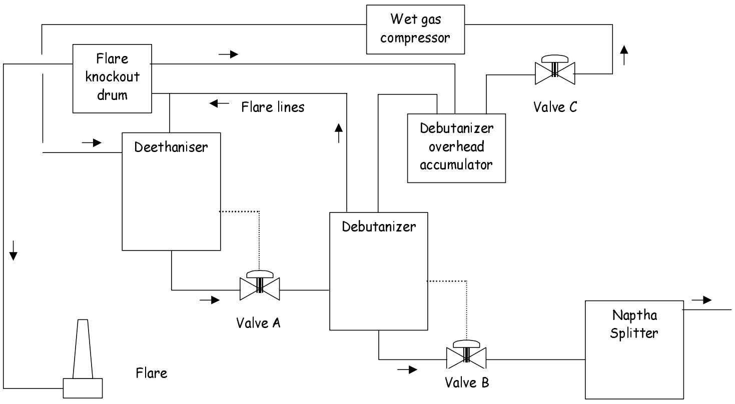

Fig. 1. High-level Overview of Components in the Fluidised Catalytic Cracking Unit

The control room displays separated crucial information that was necessary to diagnose the source of the rising pressure in the debutanizer. Rather than checking the status of valve B, the operators took action to open valve C. This allowed liquid in the full overhead accumulator to flow back into a recovery section of the plant but was insufficient to prevent the debutanizer from becoming logged with fluid entering from the deethanizer. Again, the debutanizer vented to the flare line. Opening valve C created a flow of fluid into previous 'dry' stages of the process that eventually caused a compressor trip. Large volumes of gas now had nowhere to go within the process and had to be vented to the flare stack to be burned off. At this stage, the volume of materials in the flare knockout drum was further increased by attempts to use fire hoses to drain the flooding from the dry stage directly into the flare line. However, this enabled the wet gas compressor to be restarted. This should have made matters better by increasing the flow of materials through the unit but had the unwanted effect of causing a further increase of pressure in the debutanizer. The operators responded again by opening valve C causing a further trip of the compressor. More materials were vented to an already full flare drum. Liquid was forced into a corroded discharge pipe, which broke at an elbow bend causing 20 tonnes of highly flammable hydrocarbon to be discharged. The resulting vapour cloud ignited causing damage estimated to be in excess of £50 million. This case study has been chosen to illustrate the remainder of the paper because it is typical of the way in which incidents stem from the interaction between E/E/PES-related failures, operator 'error', hardware faults and management issues. It is important to observe that both the suppliers and the operators involved in the incidents that form this case study were entirely unaware of the particular failure modes before they occurred. It is also important to emphasise that the case study cannot be characterised as software or a hardware failure. It stemmed from complex interactions between a number of system components.

Elicitation Methods: Barrier and Change Analysis

Incident reporting forms are unlikely to yield sufficient information about the causes of complex incidents and accidents. Barrier and Change analysis provide high-level frameworks for thinking about the factors that should be considered when gathering additional information. Barrier analysis stems from work in energy production (US Department of Energy, 1992). The central idea is that incidents are caused when unwanted energy flows between a source and a target. The analysis progresses by examining the barriers that might prevent a hazard from affecting the targets. Analysts must account for the reasons why each barrier actually did or might have failed to protect the target. Table 1 illustrates the output from this stage. The control loops illustrated in Figure 1 were intended to prevent the hazards associated with vessels becoming over-pressurized. The E/E/PES intervened by increasing the discharge rate if sensors detected that the level of material in the vessel had risen above an acceptable limit. There was, however, no backup system to reduce input to the vessel. A hazardous situation could arise if the input exceeded the capacity of the automated system to increase the outflow from the vessel. In other words, the E/E/PES logic was based on the assumption that output could always be increased beyond the input to each stage of the process. The meta-level point is that Barrier analysis encourages designers to look beyond the immediate triggering events that led to the mishap, to design issues rather than individual operator action.

Barrier	Reason for failure?
Control loops link level in each vessel to discharge rate.	No control over input to vessel so assumes discharge rate can always exceed input rate. This was not the case during the incident.
	No secondary control loop to monitor input rate and limit it if this exceeded output. Hence there was no backup or secondary control system.
	Key sensors provided erroneous information to control system.
Control system displays linked to multiple level alarms.	Operators could not identify reason for alarms, especially that valve B was closed, because displays were grouped according to sub-processes hence it was difficult to gain over view of the system state.
	Operators were progressively overloaded with alarms. In the last eleven minutes they were expected to read and confirm 275 individual alarms, with similar high severity levels.

Table 1. Example Barrier Analysis

Table 2 provides an example of the manner in which change analysis looks at the differences that occur between the actual events leading to an incident and 'normal' or 'ideal' operating practices. The first column describes the ideal condition or the condition prior to the incident. This is an important distinction because the causes of adverse events often stem from inappropriate practices that continue for many months. In such circumstances, the change analysis would focus less on the conditions immediately before the incident and more on the reasons why practice changed from the ideal some time before the mishap. As with Barrier analysis, this technique encourages investigators to gather information about the longer-term, less direct, factors that contribute to a mishap.

Prior/Ideal Condition	Present Condition	Effect of Change
All modifications that have safety implications should be considered by a formal hazard assessment.	The modification to the flare system that reduced the normal transfer capacity to storage tanks so that material could be recovered back to the production process was not considered in any formal safety assessment.	Assumed that operators would manually restore system configuration so that high-capacity pumps could remove excess materials from flare knockout drum. Operators were unprepared to do this both by lack of training and difficulty of diagnosing the state of their system.
Maintenance procedures should ensure that E/E/PES applications have accurate and timely information to meet control requirements.	After the incident some 40 control loops were tested. 24 required maintenance from minor mechanical damage to major faults. Of these, only the faulty debutanizer outlet valve (labeled B) occurred on the day of the incident. All of the rest were either known about or had not been detected for some time.	Maintenance issues created latent conditions for which the lightening strike acted as a catalyst. Key readings confused operators. E/E/PES sensors downstream reported that valve B was closed. However, the flow indicator closest to the valve indicated a flow and a level in the debutanizer that was below maximum even though it was full.

Table 2. Change Analysis

Event-Based Techniques: Timelines, Accident Fault Trees and Chain of Event Models

Once necessary information has been gathered about an incident, it is often used to develop some form of graphical timeline. Timelines provide arguably the simplest form of event-based analysis technique. Figure 2 provides an example for our case study. It uses a technique that was pioneered by groups within the US National Transportation Safety Board (Johnson, 2002). Events are placed on a horizontal time-line and are grouped according to the agents involved. In this case, events relating to the debutanizer and deethanizer are separated from the operator actions and so on. Such structuring mechanisms are important if analysts are not to be overwhelmed by the mass of detail that can be obtained in the aftermath of an adverse event. A number of problems affect the use of time-lines in the reconstruction and causal analysis of E/E/PES related incidents. There will often be inconsistencies and contradictory evidence for exact timings. It can also be impossible to obtain exact timings for some events. Figure 2 illustrates a point in the investigation where we know that the level alarm for the Naptha splitter occurred at some point after 08.40 but further analysis of the alarm logs is needed to determine the exact timing for this event.

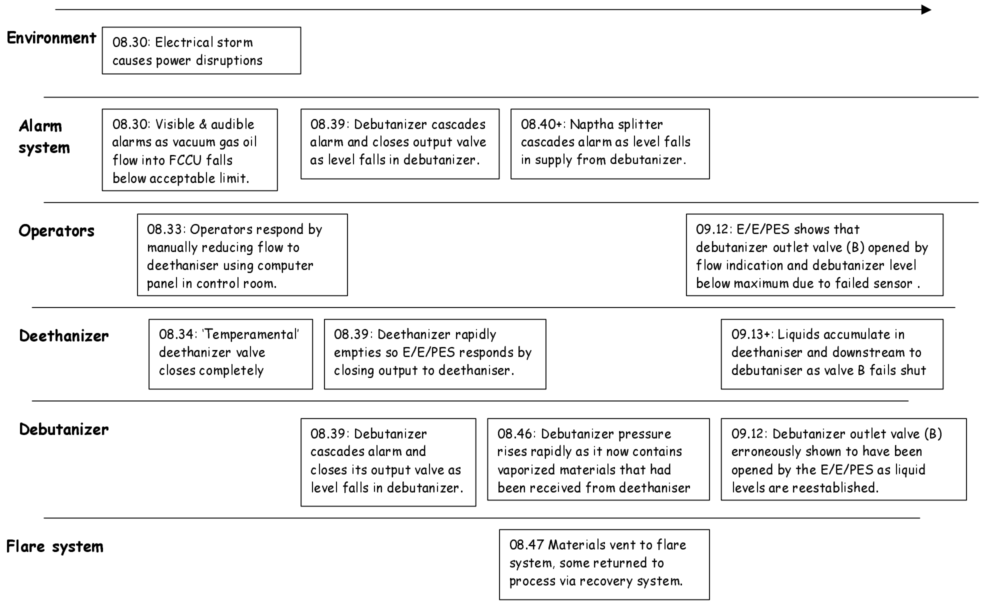
Environment
08.30: Electrical storm causes power disruptions
Alarm system
08.30: Visible & audible alarms as vacuum gas oil flow into FCCU falls below acceptable limit.
08.39: Debutanizer cascades alarm and closes output valve as level falls in debutanizer.
08.40+: Naptha splitter cascades alarm as level falls in supply from debutanizer.
Operators
08.33: Operators respond by manually reducing flow to deethaniser using computer panel in control room.
09.12: E/E/PES shows that debutanizer outlet valve (B) opened by flow indication and debutanizer level below maximum due to failed sensor .
Deethanizer
08.34: 'Temperamental' deethanizer valve closes completely
08.39: Deethanizer rapidly empties so E/E/PES responds by closing output to deethaniser.
09.13+: Liquids accumulate in deethaniser and downstream to debutaniser as valve B fails shut
Debutanizer
08.39: Debutanizer cascades alarm and closes its output valve as level falls in debutanizer.
08.46: Debutanizer pressure rises rapidly as it now contains vaporized materials that had been received from deethaniser
09.12: Debutanizer outlet valve (B) erroneously shown to have been opened by the E/E/PES as liquid levels are reestablished.
Flare system
08.47 Materials vent to flare system, some returned to process via recovery system.

Fig. 2. High-level Timeline of the Case Study Incident

A number of attempts have been made to extend fault-tree notations from the design of safety-critical systems to support the analysis of incidents and accidents. This approach has the obvious benefit that engineers who are trained in the existing use of Fault Trees can apply their knowledge and tool support to investigate the causes of adverse events. Figure 3 shows how events that contribute to a mishap are represented as rectangles. Logic gates are used to describe relationships between these events. In this case, the tree only includes 'AND' gates. For example, the bottom right sub-tree shows that the 'Excess material was recovered from the flare system at two slow a rate' as a result of a 'Decision to disable the high capacity pump to slops' AND the 'Operators failure to recognize the need to manually set-up discharge pumps' AND the lack of any 'formal risk assessment of the modifications to the flare system'.

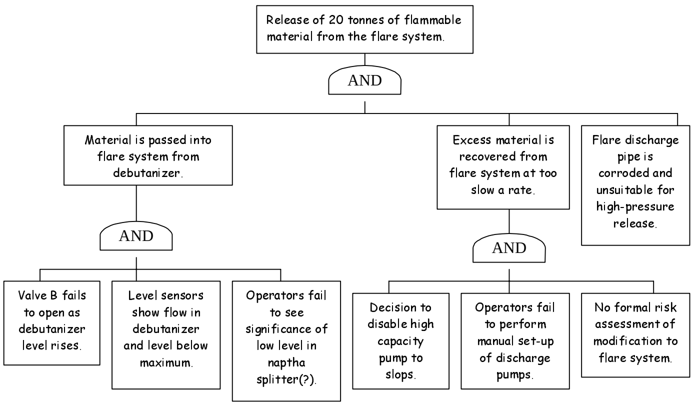
Release of 20 tonnes of flammable material from the flare system.
AND
Material is passed into flare system from debutanizer.
Excess material is recovered from flare system at too slow a rate.
Flare discharge pipe is corroded and unsuitable for high-pressure release.
AND
AND
Valve B fails to open as debutanizer level rises.
Level sensors show flow in debutanizer and level below maximum.
Operators fail to see significance of low level in naptha splitter(?).
Decision to disable high capacity pump to slops.
Operators fail to perform manual set-up of discharge pumps.
No formal risk assessment of modification to flare system.

Fig. 3. Overview of an Accident Fault Tree

Figure 3 illustrates a number of important differences that distinguish the use of fault trees in accident investigation from their more conventional use to support the design of safety-critical systems. OR gates are not used. This would imply uncertainty in the reconstruction – there would be two alternative paths of events leading to the failure. Such uncertainty is, in general, avoided in incident investigation unless analysts are explicitly looking for alternative failure mechanisms that might lead to slightly different mishaps in the future. There are further differences between accident fault trees and the use of this technique for design. For example, it is unclear how to represent the events that occur in the immediate aftermath of a mishap. This is important because the response to an incident can help to determine the eventual outcome. In conventional fault-trees the analysis stops with a potential hazard. This is particularly significant in our case study. There were no plans to deal with a large fire burning for longer than twenty-four hours. This placed extreme demands on local water supplies that were needed both to fight the fire and to cool nearby vessels that might have been affected by rising temperatures.

There are several more complex techniques for plotting out the events that contribute to accidents and incidents. Figure 4 illustrates a Failure Event Tree that is similar to the output from Events and Causal Factors charting (ECF), Multilinear Events Sequencing (MES) and Sequential Timed Event Plotting (STEP) (US Department of Energy, 1992). A sequence of events leads to the mishap. These are denoted by the rectangles on the top half of the image. Outcomes are denoted by bold rectangles with dotted borders. Figure 4 also captures direct factors that influence the course of the incident but which cannot conveniently be represented by discrete events. These are denoted by rectangles with a double line border, such as ‘No second back-up feedback control loop to ensure input flow reduced or shut-off when material accumulates’. Finally, Figure 4 captures a series of less direct factors that contribute to the incident. Many would argue that these factors represent the root causes of an accident or near miss. They include observations that there were ‘poor maintenance procedures’ and ‘alarms cascade with low prioritization and requirement for explicit acknowledgement from operators’. We have extended the basic form of Failure Event Trees by shading those events that directly refer to intervention by E/E/PES related systems. This illustrates the way in which programmable devices compound operator ‘error’, maintenance failures and many other types of events during the course of most mishaps. They are seldom the ‘only cause’ of adverse events in complex, safety-critical systems.

Many event based techniques, including Failure Event Trees, exploit counter-factual arguments to distinguish root causes from less significant events. These arguments take the form ‘if X did not occur then the accident/incident would have been avoided’. This form of argument is ‘counterfactual’ because we know that the accident or incident did take place. We are trying to imagine ways in which we might have avoided the failure. Analysts use this form of reasoning by looking at the event closest to the incident. In Figure 4, we ask would the mishap still have occurred if liquid had not been forced from the full flare drum into the corroded discharge pipe. If the answer is yes and the mishap would still have happened then this event cannot be a candidate root cause of the incident. If the answer is no and the mishap would not have occurred without this event then we can argue that it was necessary for the incident to occur so it can be considered as a root cause. The process continues for each of the mishap events shown in the diagram. In this example, it can reasonably be argued that the incident would still have occurred even if the pipe had not failed. The increasing pressure in the flare drum is likely to have led to another form of failure. The search for root causes then continues amongst the previous events. This form of reasoning might, for instance, be used to focus on the operators’ decision to open valve C rather than examine the possibility that their control system was presenting misleading information about the state of valve B.

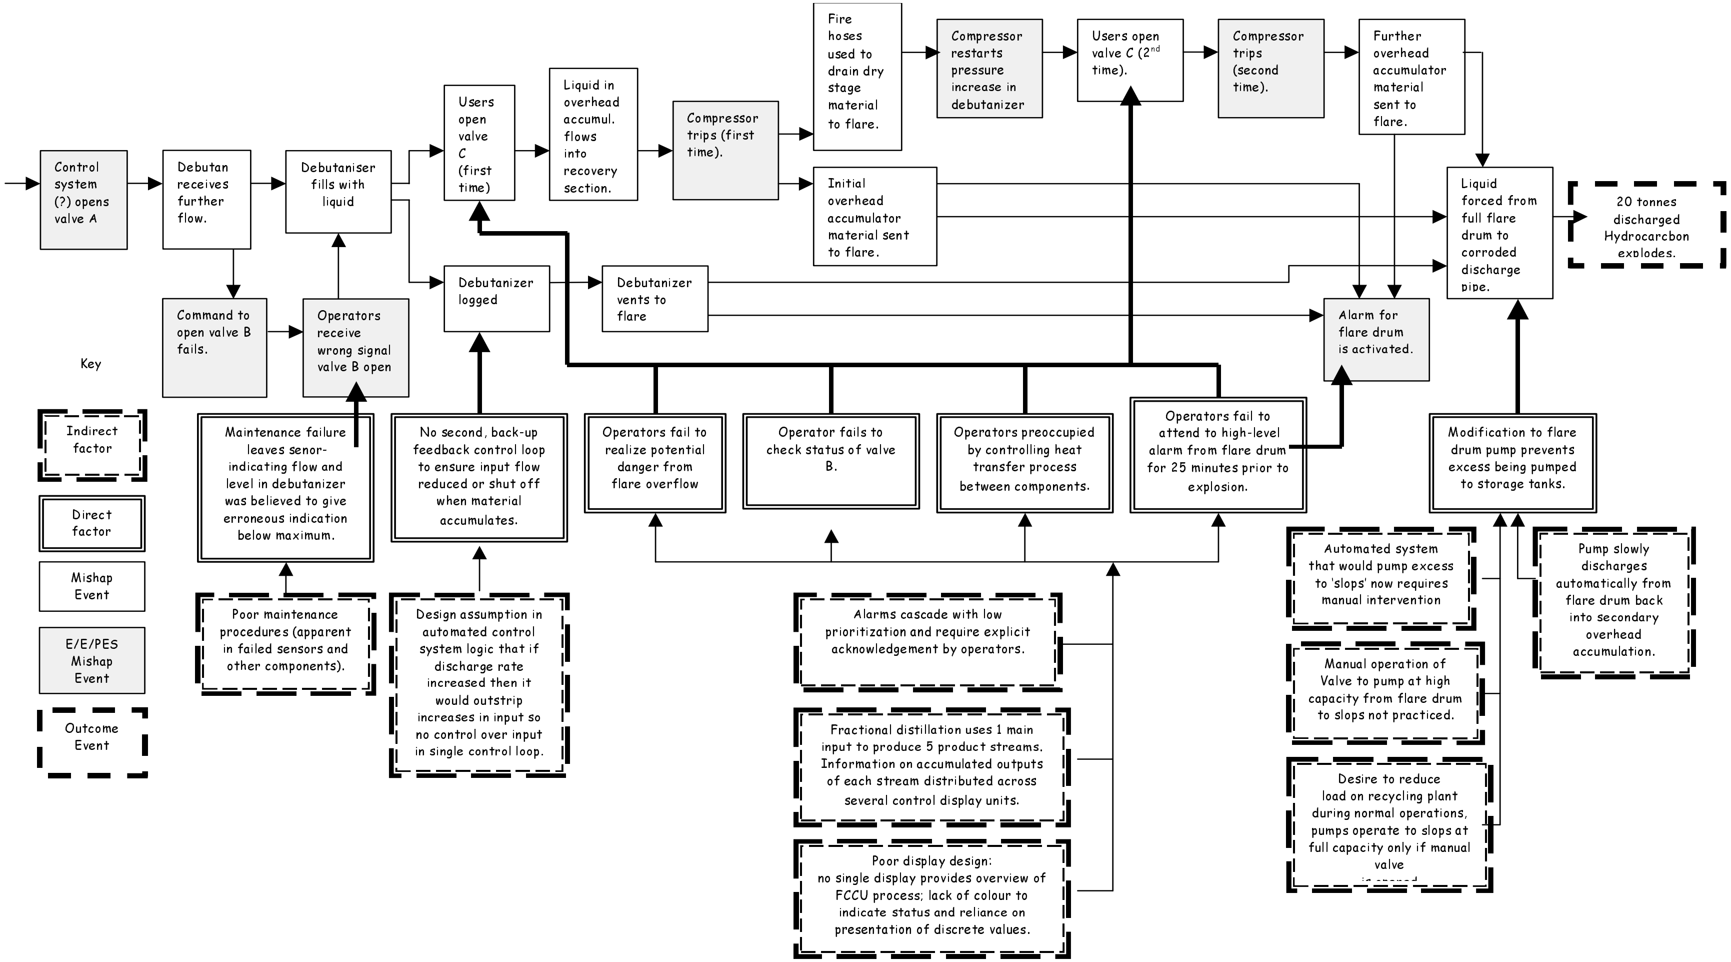
Control system (?) opens valve A
Debutan receives further flow.
Debutaniser fills with liquid
Command to open valve B fails.
Operators receive wrong signal valve B open
Users open valve C (first time)
Debutanizer logged
Liquid in overhead accumul. flows into recovery section.
Debutanizer vents to flare
Compressor trips (first time).
Fire hoses used to drain dry stage material to flare.
Initial overhead accumulator material sent to flare.
Compressor restarts pressure increase in debutanizer
Users open valve C (2nd time).
Compressor trips (second time).
Further overhead accumulator material sent to flare.
Liquid forced from full flare drum to corroded discharge pipe.
20 tonnes discharged Hydrocarcbon explodes.
Alarm for flare drum is activated.
Key
Indirect factor
Direct factor
Mishap Event
E/E/PES Mishap Event
Outcome Event
Maintenance failure leaves senor-indicating flow and level in debutanizer was believed to give erroneous indication below maximum.
No second, back-up feedback control loop to ensure input flow reduced or shut off when material accumulates.
Operators fail to realize potential danger from flare overflow
Operator fails to check status of valve B.
Operators preoccupied by controlling heat transfer process between components.
Operators fail to attend to high-level alarm from flare drum for 25 minutes prior to explosion.
Modification to flare drum pump prevents excess being pumped to storage tanks.
Poor maintenance procedures (apparent in failed sensors and other components).
Design assumption in automated control system logic that if discharge rate increased then it would outstrip increases in input so no control over input in single control loop.
Alarms cascade with low prioritization and require explicit acknowledgement by operators.
Fractional distillation uses 1 main input to produce 5 product streams. Information on accumulated outputs of each stream distributed across several control display units.
Poor display design: no single display provides overview of FCCU process; lack of colour to indicate status and reliance on presentation of discrete values.
Automated system that would pump excess to 'slops' now requires manual intervention
Manual operation of Valve to pump at high capacity from flare drum to slops not practiced.
Desire to reduce load on recycling plant during normal operations, pumps operate to slops at full capacity only if manual valve is opened
Pump slowly discharges automatically from flare drum back into secondary overhead accumulation.

Fig. 4. A Failure Event Tree

Flow Charts and Taxonomies: MORT and PRISMA

Management Oversight and Risk Trees (MORT) provide the best-known example of a flow charting approach to the identification of causal factors (W. Johnson, 1980). Figure 5 provides an abbreviated version of a MORT diagram. Investigators first consider the top levels of the tree. They must ask themselves whether the mishap was the result of an omission of some management function and whether the incident occurred from a risk that had already been recognized. In the tree, the term LTA refers to a 'less than adequate' performance of some necessary activity. If there was an oversight problem then analysis progresses to the next level of the tree. Investigators are encouraged to consider both what happened and why it happened. The reasons why an oversight might occur include less than adequate management policy, implementation or risk assessment. Investigators work their way through the tree shown in Figure 5 until they reach a terminal node that describes the incident under consideration. These terminal nodes are not shown here. The full MORT diagram contains several hundred components. However, these leaves describe the detailed managerial causes of adverse events. For example, the analysis of the case study might begin by asking whether the oversight during development was adequate. If it was not then we can begin to analyze what happened during the incident by going down the far left branch of the tree. After having identified what occurred, analysts consider the right branches including the reasons why management might have been less than adequate. The right most sub-branch encourages analysts to determine whether this was due to incorrect goals, to problems in the technical information systems that were available to management, to inadequate hazard analysis or problems in the safety program review process. For instance, modifications to the E/E/PES controlled high-capacity flare excess pumping system could be a result of inadequate hazard analysis because the danger of operators failing to manually reconfigure the lower capacity reclamation pumps was not considered in sufficient detail.

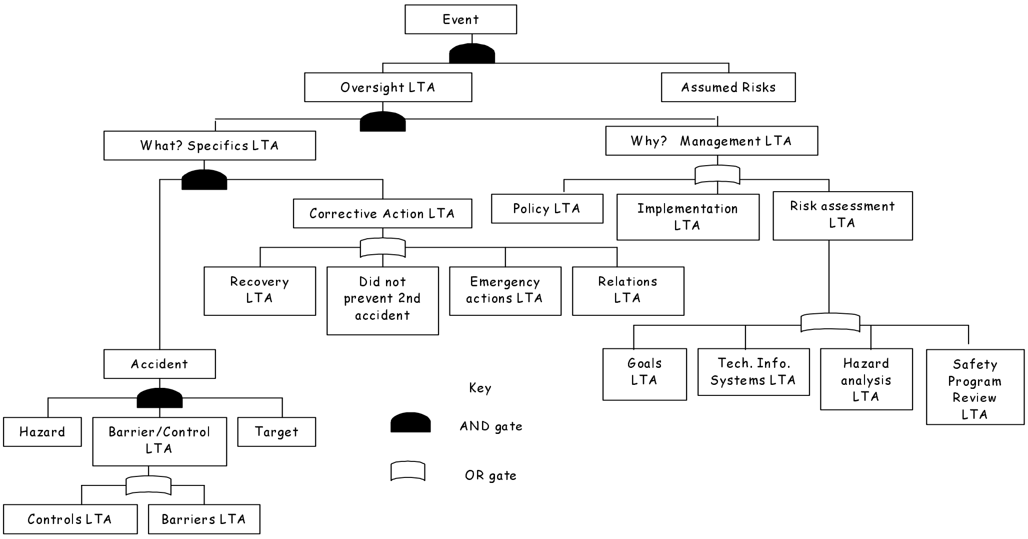

Fig. 5. Abbreviated form of a MORT diagram

PRISMA provides a further example of a flow chart technique that can be used to identify the causes of E/E/PES related incidents. Unlike MORT, it starts with an initial reconstruction based on an accident fault tree (van der Schaaf, 1992). The leaf or terminal nodes on the tree are then classified to identify more generic causes using a flow chart. Figure 6 illustrates a PRISMA flow chart that was developed to identify higher-level causal factors in the process industries. As can be seen, each terminal node is associated with a particular abbreviation such as TE for a technical, engineering related cause. The developers of the PRISMA approach encourage investigators to extend the classification to support their particular domain of interest. For example, medical versions include 'patient related factors' as a potential cause in healthcare incidents. In our case study, we might extend the flow chart to explicitly consider more detailed technical factors than those shown in Figure 6. For instance, we might introduce nodes to capture failures that are due to the discrepancies between the state of process components and their representation on the control system display. Similarly, the flowchart might be refined to help investigators categorise incidents in which E/E/PES embodied hazardous assumptions about process components. This would include the erroneous use of the argument that a single control loop would suffice because it would always be possible to increase sub-processes' outflow beyond their input.

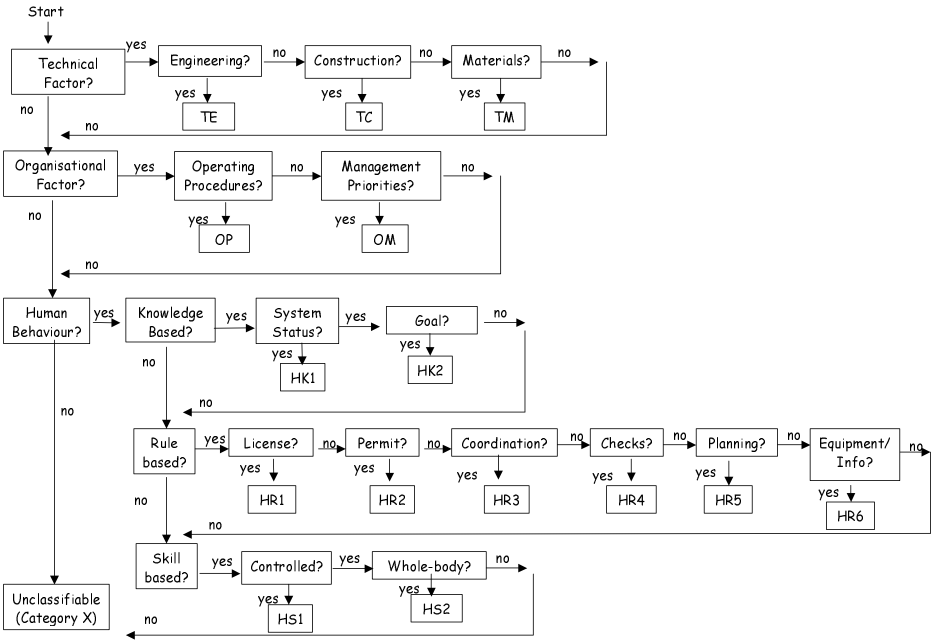

Fig. 6. PRISMA Flow Chart (van der Schaaf, 1992, 1996)

An important strength of flow-chart methods such as PRISMA is that the generic causal classification can direct investigators towards a number of general solutions. Table 6 illustrates a classification action matrix. This shows that if, for example, an incident were due to problems with management priorities then subsequent recommendations might focus more on 'bottom-up communication'. If incidents continue to recur with the same set of causal factors then safety managers might decide that the remedies advocated in Table 6 are ineffective and should be revised. In our case study, it might be advocated that modifications be made to introduce secondary control loops for the E/E/PES monitoring levels in each of the process stages. Such a detailed remedial action could only be represented in a classification/action matrix if the associated flow chart were extended to a similar level of complexity.

	External Factors (O-EX)	Knowledge Transfer (OK)	Operating procedures (OP)	Manag. priorities (OM)	Culture (OC)
Inter-departmental communication	X				
Training and coaching		X			
Procedures and protocols			X		
Bottom-up communication				X	
Maximise reflexivity					X

Table 1. Example PRISMA Classification/Action Matrix Van Vuuren (1998)

Accident Models: TRIPOD and STAMP

A number of causal analysis techniques have been developed around ‘accident models’; these provide templates that investigators must instantiate with the details of a particular mishap. The Tripod approach builds on the notion that most adverse events and near misses are caused by more general failure types: Hardware; Maintenance management; Design; Operating procedures; Error-enforcing conditions; Housekeeping; Incompatible goals; Communication; Organisation; Training; Defence planning. Software is a notable omission from this list and must certainly be included. Figure 7 illustrates a Tripod graphical model that can be used to show how specific instances of these general failure types combine to create an incident or accident. Elements of barrier analysis are used to show to associate a number of active failures with each of the defences that did not protect the target. These active failures can be thought of as the immediate events leading to the incident. The context in which they can occur is often created by a number of preconditions. For instance, the preconditions for the accumulation of material in the debutanizer were valve B sticking shut while the E/E/PES sensors continued to detect both a flow and a level below the maximum. These active failures and preconditions were, in turn, due to latent problems in the maintenance procedures that ensured E/E/PES functionality.

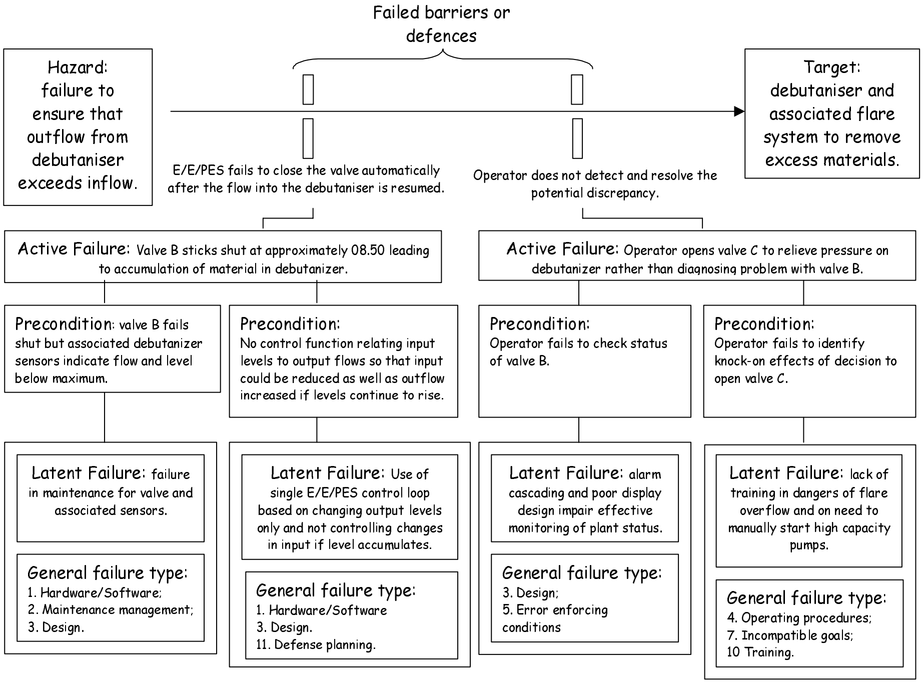
Failed barriers or defences
Hazard: failure to ensure that outflow from debutaniser exceeds inflow.
Target: debutaniser and associated flare system to remove excess materials.
E/E/PES fails to close the valve automatically after the flow into the debutaniser is resumed.
Operator does not detect and resolve the potential discrepancy.
Active Failure: Valve B sticks shut at approximately 08.50 leading to accumulation of material in debutanizer.
Active Failure: Operator opens valve C to relieve pressure on debutanizer rather than diagnosing problem with valve B.
Precondition: valve B fails shut but associated debutanizer sensors indicate flow and level below maximum.
Precondition: No control function relating input levels to output flows so that input could be reduced as well as outflow increased if levels continue to rise.
Precondition: Operator fails to check status of valve B.
Precondition: Operator fails to identify knock-on effects of decision to open valve C.
Latent Failure: failure in maintenance for valve and associated sensors.
General failure type: 1. Hardware/Software; 2. Maintenance management; 3. Design.
Latent Failure: Use of single E/E/PES control loop based on changing output levels only and not controlling changes in input if level accumulates.
General failure type: 1. Hardware/Software 3. Design. 11. Defense planning.
Latent Failure: alarm cascading and poor display design impair effective monitoring of plant status.
General failure type: 3. Design; 5. Error enforcing conditions
Latent Failure: lack of training in dangers of flare overflow and on need to manually start high capacity pumps.
General failure type: 4. Operating procedures; 7. Incompatible goals; 10 Training.

Fig. 7. Example Application of a TRIPOD General Failure Types

Leveson’s Systems Theory Accident Modeling and Process (STAMP) exploits elements of control theory to help identify causal factors. This is motivated by the observation that mishaps occur when external disturbances are inadequately controlled. Control failures can arise from ‘dysfunctional interactions’ between system components. For example, if one subsystem embodies inappropriate assumptions about the performance characteristics of another process component. In this view, mishaps do not stem from events but from inappropriate or inadequate constraints on the interactions among the elements that form complex, safety-critical applications. Safety is viewed as a dynamic property of the system because the degree to which a system satisfies those constraints will continually evolve over time. Figure 8 illustrates this approach. Arrows represent communication and control flows. Rectangles are entities, including people, systems and organizations. STAMP control analysis extends from the operator, and the systems under their immediate control to also consider the relationships between project and company management, between management and regulatory agencies and between regulation and system vendors. After having conducted this extended form of control analysis, STAMP considers each of the control loops that are identified in the ‘socio-technical system’. Potential mishaps stem from missing or inadequate constraints on processes or from the inadequate enforcement of a constraint that contributed to its violation. Table 2 illustrates the general classification scheme that guides this form of analysis. Analysis progresses by examining each of the arrows in the control model to see whether any of the flaws in Table 2 can be identified in the relationships that they represent. It might be argued that there were unidentified hazards in the control loops between the control system, valve B and the debutaniser. Similarly, subsequent investigation might identify flaws in the creation process that led to the operators' control system display of the state of process components.

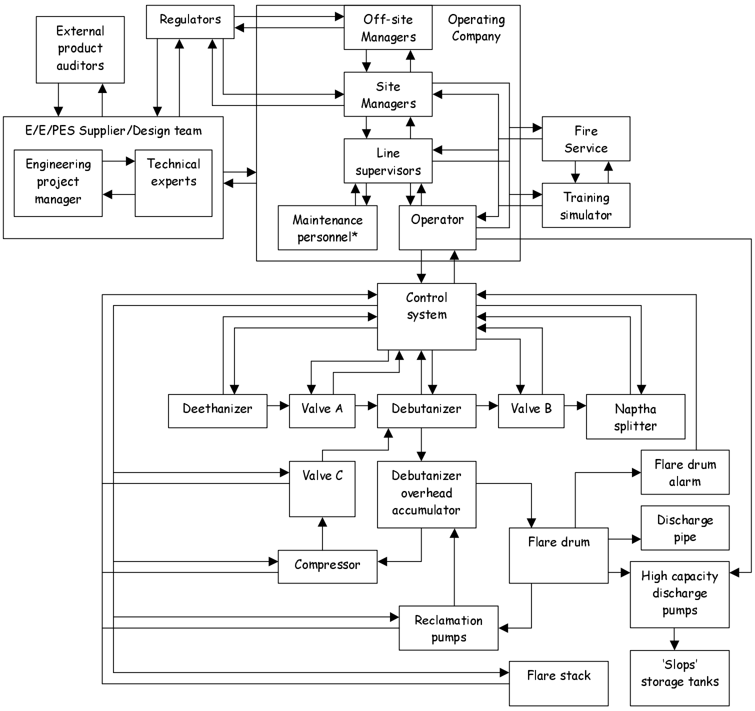

Fig. 8. Example Control Model from STAMP

1. Inadequate Enforcements of Constraints (Control Actions)

1.1 Unidentified hazards

1.2 Inappropriate, ineffective or missing control actions for identified hazards

1.2.1 Design of control algorithm (process) does not enforce constraints

- Flaws in creation process
- Process changes without appropriate change in control algorithm (asynchronous evolution)
- Incorrect modification or adaptation.

1.2.2 Process models inconsistent, incomplete or incorrect (lack of linkup)

- Flaws in creation process
- Flaws in updating process (asynchronous evolution)
- Time lags and measurement inaccuracies not accounted for

1.2.3 Inadequate coordination among controllers and decision makers

2 Inadequate Execution of Control Action

2.1 Communication flaw

2.2 Inadequate actuator operation

2.3 Time lag

3. Inadequate or Missing Feedback

3.1 Not provided in system design

3.2 Communication flow

3.3 Time lag

3.4 Inadequate sensor operation (incorrect or no information provided)

Table 2. Control Flaws leading to Hazards (Leveson, 2002)

Argumentation Techniques: WBA and CAE

Several techniques have been developed to help ensure that investigators form ‘reasonable’ causal arguments from the evidence that is embodied in timelines and other reconstructions. Ladkin and Loer’s (1998) Why-Because Analysis uses a graphical notation to reconstruct sequences of events leading to a mishap. The angled arrows shown in Figure 9 illustrate this. As can be seen, the E/E/PES opens valve C before it issues the unsuccessful command to open valve B. All of this occurs before the debutanizer fills with material. It is important to stress, however, that this sequential information does not imply causation. Mathematically based proof techniques provide a method for establishing such causal arguments. Informally, we must demonstrate that we have identified sufficient causes for an ‘effect’ to occur. Once analysts are convinced that they have considered a sufficient set of causal factors for an effect they can then revise the WBA diagram illustrated in Figure 9. A double arrow denotes causal relationships =>>. As can be seen, the overhead accumulator material sent to the flare is sufficient in this diagram for the flare drum alarm to be activated. This transition from temporal sequences to more rigid causal relationships can produce insights that are not apparent in purely event-based approaches, such as timelines. For example, we might consider that the operators’ initial action to open valve C made them more likely to repeating this intervention when faced with another warning of increasing pressure in the debutanizer. This is explicitly represented in the WBA diagram of Figure 9 but was not previously included in the event-driven approach of Figure 4’s Failure-Event tree. The most striking feature of WBA is that it provides a set of mathematically based procedures that analysts must follow in order to replace the angled arrows of a temporal sequence with the double headed arrows of the causal relationship. These procedures are necessary to ensure that we have established sufficient causes for the effect to occur. They are based on arguments of the form ‘A causes B’ if B is true in possible worlds that are close to those in which A is true, which can in turn be given a counterfactual semantics. Ladkin and Loer also provide a range of additional proof rules that can be used to ensure both the consistency and sufficiency of arguments about the causes of a mishap.

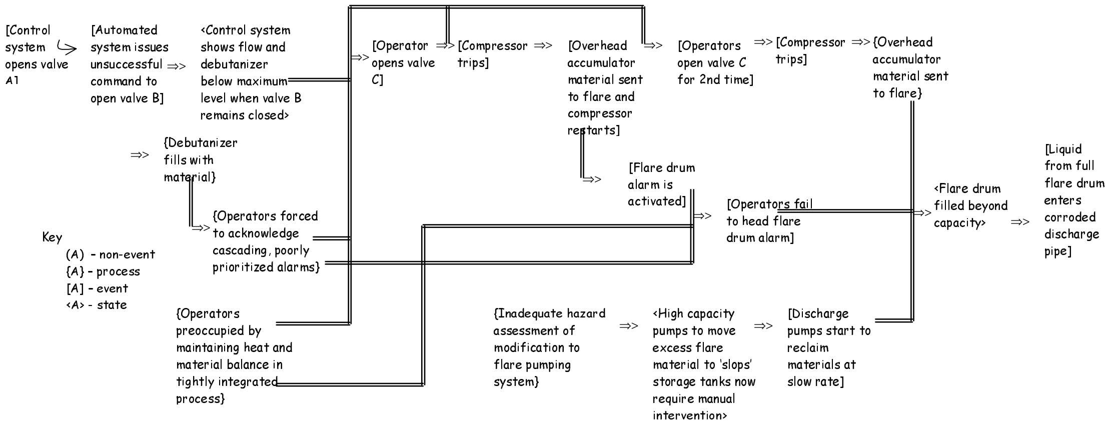
[Control system opens valve A1
[Automated system issues unsuccessful command to open valve B]
<Control system shows flow and debutanizer below maximum level when valve B remains closed>
[Operator opens valve C]
[Compressor trips]
[Overhead accumulator material sent to flare and compressor restarts]
[Operators open valve C for 2nd time]
[Compressor trips]
{Overhead accumulator material sent to flare}
{Debutanizer fills with material}
[Flare drum alarm is activated]
[Liquid from full flare drum enters corroded discharge pipe]
<Flare drum filled beyond capacity>
[Operators fail to head flare drum alarm]
{Operators forced to acknowledge cascading, poorly prioritized alarms}
Key
(A) - non-event
{A} - process
[A] - event
<A> - state
{Operators preoccupied by maintaining heat and material balance in tightly integrated process}
{Inadequate hazard assessment of modification to flare pumping system}
<High capacity pumps to move excess flare material to 'slops' storage tanks now require manual intervention>
[Discharge pumps start to reclaim materials at slow rate]
=>>

Fig. 9. Example WBA Diagram

Conclusion, Analysis and Evidence (CAE) diagrams help designers to map out the competing arguments for and against particular conclusions or recommendations in the aftermath of a complex incident. This approach lacks the consistency and completeness checks that are provided by the formal reasoning in the WBA technique. However, the reliance on a graphical formalism together with less strict rules on how to conduct the analysis can arguably reduce the costs and increase the flexibility of this approach. Figure 10 provides an example of a CAE diagram. Rectangles are connected to form a network that summaries arguments about an incident or accident. As the CAE name suggests, the rectangles labeled with a C are used to denote conclusions or recommendations, those labeled with an A are lines of analysis while the E rectangles denote evidence. Lines are drawn to show those lines of analysis that support particular conclusions. For example, the recommendation that the operators’ safety management system be revised to explicitly store, retrieve and review incident information from other plants (C.1) is taken directly from the primary recommendation of the official report into this incident. The conclusion is supported by the observation that previous incidents were caused by a similar failure to assess the hazards of process modifications (A1.1). The evidence for this assertion is provided by the Grangemouth hydrocracker incident (E.1.1.1) and by the Flixborough explosion (E.1.1.2). It is important to note that Figure 10 also captures contradictory arguments. For instance, the dotted line in the first network denotes that the existing safety management system is not guaranteed to have acted on previous incident information even if it had been gathered more explicitly (A.1.2) given that there were other failings in the monitoring of maintenance and modification tasks (E1.2.1). As can be seen from Figure 10, CAE diagrams capture general arguments about incidents and accidents. For example, a conclusion might refer to a recommended action; it need not simply capture a causal relationship. It is also important to mention that this technique was specifically developed to enable investigators to sketch out the arguments that might appear in an incident report. This helps to ensure that any document avoids contradictory arguments.

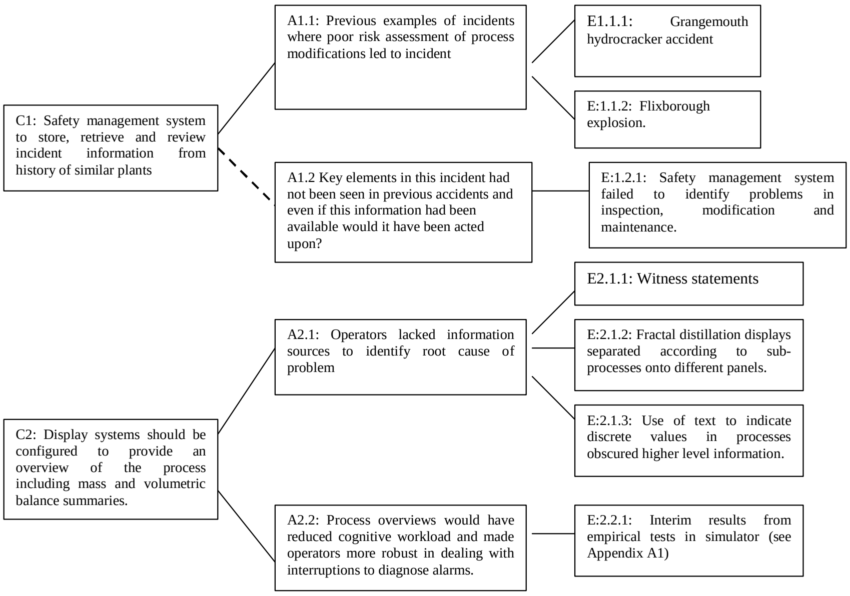
A1.1: Previous examples of incidents where poor risk assessment of process modifications led to incident
E1.1.1: Grangemouth hydrocracker accident
E:1.1.2: Flixborough explosion.
C1: Safety management system to store, retrieve and review incident information from history of similar plants
A1.2 Key elements in this incident had not been seen in previous accidents and even if this information had been available would it have been acted upon?
E:1.2.1: Safety management system failed to identify problems in inspection, modification and maintenance.
E2.1.1: Witness statements
A2.1: Operators lacked information sources to identify root cause of problem
E:2.1.2: Fractal distillation displays separated according to sub-processes onto different panels.
E:2.1.3: Use of text to indicate discrete values in processes obscured higher level information.
C2: Display systems should be configured to provide an overview of the process including mass and volumetric balance summaries.
A2.2: Process overviews would have reduced cognitive workload and made operators more robust in dealing with interruptions to diagnose alarms.
E:2.2.1: Interim results from empirical tests in simulator (see Appendix A1)

Fig. 10. Example of a CAE Diagram

Comparisons

Most companies and regulatory organizations lack the resources to train investigators in a range of different causal analysis techniques. It is, therefore, important to help managers focus finite resources by identifying those techniques that are best suited to analyzing the causes of computer-related incidents. Table 3, therefore, presents a subjective assessment of whether each of the previous approaches can be used to uncover problems in particular stages of the E/E/PES lifecycle or in requirements that must be satisfied across those different stages, such as staff competency. Lifecycle components and common requirements were derived from an analysis of the IEC 61508 standard. This decision was partly justified by pragmatics; this approach was recommended in consultation with the UK Health and Safety Executive. Other lifecycle models could have been used. The initial assessments in Table 3 were validated during consultations with members of the HSE, with experts on the IEC 61508 standard and by the members of several safety-critical software consultancies. We are currently extending this exercise to incorporate the opinions of safety managers across the UK process industries.

	Elicitation and Analysis techniques		Event Based Techniques			Flowcharts and taxonomies		Accident Models		Argumentation Techniques	
	Barrier Analysis	Change Analysis	Timelines	Accident Fault Trees	ECF	MORT	PRISMA	TRIPOD	STAMP	WBA	CAE
IEC 61508 Lifecycle phase											
Concept	F	F	U	U	P	F	P	F	P	U	F
Overall Scope	F	F	U	U	P	F	P	F	P	U	F
Hazard & Risk Assessment	P	P	P	P	P	F	P	P	F	U	F
Overall Safety Requirements	F	F	U	U	P	P	P	F	F	U	F
Allocation	F	P	P	U	F	P	P	F	P	U	U
Planning of Validaton, Operation & maintenance	U	P	P	P	F	F	F	U	P	P	U
Realisation	U	F	F	P	F	U	P	U	F	F	U
Installation & Commissioning	U	P	F	P	F	P	P	P	P	F	P
Validation	P	P	F	P	F	P	P	P	U	F	P
Operation & Maintenance	P	F	F	P	F	P	P	F	F	F	P
Modification	U	F	F	P	F	P	P	U	F	F	P
IEC 61508 Common Requirements											
Competency	P	P	P	P	F	P	P	F	P	P	P
Lifecycle	U	P	P	P	F	P	P	P	P	P	P
Verification	P	P	P	P	F	P	F	P	P	P	P
Safety management	P	P	P	P	F	P	P	P	P	P	P
Documentation	P	P	P	P	F	P	P	P	P	P	P
Functional safety assessment	P	P	P	P	F	P	P	P	P	P	P

Key: (U)nsupported, (P)artially supported, (F)ully supported

Table 3. Degree of Support for Mapping from Products of Causal Analysis Technique to Failures in IEC 61508 Lifecycle Phases and Common Requirements

It is important to provide a brief rationale for some of the assessments captured in table 3. Barrier and Change analysis are similar because they provide a means of investigating incidents at a relatively high level of granularity. For instance, Barrier Analysis focuses on the differences between what actually did happen and what was supposed to happen. The intended behavior of the E/E/PES can partially be derived from the documentation associated with IEC 61508 development and by other legal and regulatory documents. Timelines differ from these approaches because they model what actually happened during an E/E/PES related incident. Problems arise when IEC 61508 requirements cannot be directly related to particular events. For example, the common requirement to ensure competency would have to be represented as specific events that were intended to ensure this requirement. The lack of temporal information in accident fault trees creates problems for investigators who must reason the detailed sequence of operations executed by an E/E/PES in the course of an adverse event. This also creates problems for the analysis of IEC 61508 requirements. There are dependencies between the various stages of the development lifecycle. Hazard and risk assessment should precede operation and maintenance. It is difficult to represent the violation of such requirements using this modeling and analysis technique.

MORT supports the identification of problems in the management mechanisms that are intended to protect safety-critical systems (Johnson, 1980). Table 3 therefore denotes that this causal analysis technique might support the analysis of failures in the IEC 61508 common requirements relating to ensuring competence, establishing safety management procedures etc. The subjective assessment of PRISMA in Table 3 is less a reflection of the underlying ideas than an assessment of the existing classification schemes. For example, the PRISMA flow chart illustrated in previous sections considers engineering, construction and materials as key issues in the technical reasons for an adverse event. We could extend this flowchart to represent the requirements associated with realization phase in Table 3. The initial list of TRIPOD general failure types did not include software failure. As with MORT and PRISMA, however, this could be rectified through the subsequent tailoring of the approach to support the analysis of E/E/PES related incidents. The STAMP constraint checklist guides the identification of problems in risk and hazard assessment. It offers less support for identifying the meta-level validation processes that must be used to ensure that constraints between control entities are satisfied. Introducing more sophisticated hierarchical control models could do this. WBA provides rules to establish that a causal argument is correct without predetermining what types of failures or behaviors that argument is about. This makes it difficult to classify the degree of support that this approach provides as a tool to identify the failure of IEC 61508 requirements. This flexibility is achieved at a cost in terms of the level of skill and expertise that must be acquired before the technique can be applied. CAE lacks the formal guidance of WBA but again is not specifically tailored for E/E/PES related incidents. Hence, it can be difficult for investigators to demonstrate that their application of the approach captures all necessary causal information.

Conclusions

This paper has provided a brief overview of causal analysis techniques for Electrical, Electronic or Programmable, Electronic Systems (E/E/PES) related incidents. We have identified several main classes: Elicitation and Analysis Techniques, such as Barrier Analysis; Event-based techniques, including Accident fault trees; Flow Charts, including those within the PRISMA approach; Accident Models, including the control theory model in STAMP; Argumentation Techniques, such as the counterfactual approach in WBA. The techniques differ according to the amount of investment, in terms of training and investigators' time, that is required in order to apply them. They also differ radically in the degree of support that they provide in terms of the consistency that might be achieved between individuals applying the same approach to the same incident. A more detailed introduction to various causal analysis techniques and a cost-benefit survey of the various approaches can be found in Johnson (2003). The intention has, however, been to provide a basic road map for the range of approaches that might be used to analyse the causes of E/E/PES related incidents.

The closing sections of this report have presented a subjective comparison of the approaches that we have introduced. This assessment has been structured in terms of the perceived support that each technique provides for investigators who must identify failures in particular phases of the software development lifecycle. We have identified whether each technique provides full, partial or no support for identifying problems in the tasks that comprise the IEC 61508 development model. Our use of the standard was partly justified by pragmatics; other models might have been used. The widespread application of this standard has, however, simplified the validation of table 3. This validation has included consultations with several safety-critical software engineering consultancies and industry regulators. We are now engaged in a second stage involving safety managers in companies throughout the UK process industries.

Acknowledgements

Thanks are due to Bill Black (Black Safe Consulting), Mark Bowell (UK HSE) and Peter Bishop (Adelard) for providing comments on the initial draft of this document.

References

Department of Energy, DOE Guideline Root Cause Analysis Guidance Document, Office of Nuclear Energy and Office of Nuclear Safety Policy and Standards, U.S. Department of Energy, Washington DC, USA, DOE-NE-STD-1004-92, http://tis.eh.doe.gov/techstds/standard/nst1004/nst1004.pdf, 1992.

J.A. Doran and G.C. van der Graaf, Tripod-Beta: Incident Investigation and Analysis, Proceedings of the International Conference on Health, Safety and the Environment, Society of Petroleum Engineers, New Orleans, USA, 9-12 June, 1996.

P. Hudson and J. Reason and W. Wagenaar and P. Bentley and M. Primrose and J. Visser, Tripod-Delta: Pro-active Approach to Enhanced Safety, Journal of Petroleum Technology, 40, 58-62, 1994.

W.G. Johnson, MORT Safety Assurance Systems, Marcel Dekker, New York, USA, 1980.

C.W. Johnson (2003 in press), A Handbook for the Reporting of Incidents and Accidents, Springer Verlag, London, UK.

P. Ladkin and K. Loer (1998), Why-Because Analysis: Formal Reasoning About Incidents, Bielefeld, Germany, Document RVS-Bk-98-01, Technischen Fakultat der Universitat Bielefeld, Germany.

N. Leveson, (2002), A Systems Model of Accidents. In J.H. Wiggins and S. Thomason (eds) Proceedings of the 20th International System Safety Conference, 476-486, International Systems Safety Society, Unionville, USA.

T.W. van der Schaaf, Near Miss Reporting in the Chemical Process Industry, Technical University of Eindhoven, Eindhoven, The Netherlands, 1992.