A Systems Approach to Accident Analysis

Nancy Leveson, Margaret Stringfellow, and John Thomas (MIT)

1. Introduction

In this research, we investigated how applying systems thinking can be used to improve learning from events and hopefully reduce significantly the number of incidents and accidents. Systems thinking focuses more on the system rather than on human actions in order to learn how to redesign the system to reduce losses—where the system includes engineering design, construction, operations, management, and organizational structure.

In this final report, we summarize current approaches to accident/incident analysis in a petroleum industry company we will call Company X[1] and the process industry as a whole and the limitations to current practice. Then we describe what a systems thinking approach would entail. To determine whether such an approach is practical for the petrochemical industry,[2] we apply it to the Tank Overflow incident at a Company X refinery a few years ago. Finally, we compare the results of the new systems approach to accident analysis with the actual Company X refinery accident report generated using their standard accident analysis procedures..

2. Limitations in Current Approaches to Learning from Events

The classic approaches to accident investigation stem from a model of accident causation by Heinrich published in 1931 for occupational safety. Heinrich's Domino Model was very influential in shifting the focus in occupational safety from unsafe conditions in plants to human error. He compared the general sequence of accidents to five dominoes standing on end in a line (see Figure 1). When the first domino falls, it automatically knocks down its neighbor and so on until the injury occurs. In any accident sequence, according to this model, ancestry or social environment leads to a fault of a person, which is the proximate cause for an unsafe act or condition (mechanical or physical), which results in an accident, which leads to an injury.

The basic Domino Model proved to be inadequate for complex systems and other models were developed, but some of the assumptions of the Domino Model persist such as there being a single or "root cause" of an accident and the idea of dominos or chains of events, each directly causing or leading to the next one in the chain. It also lives on in the emphasis on human error in identifying accident causes.

[1] The accident is a real one but we have omitted the name of the company to protect their privacy.

[2] The approach has been applied effectively in accident analysis and other aspects of safety in a variety of industries including commercial aviation, defense, railroads, and pharmaceuticals.

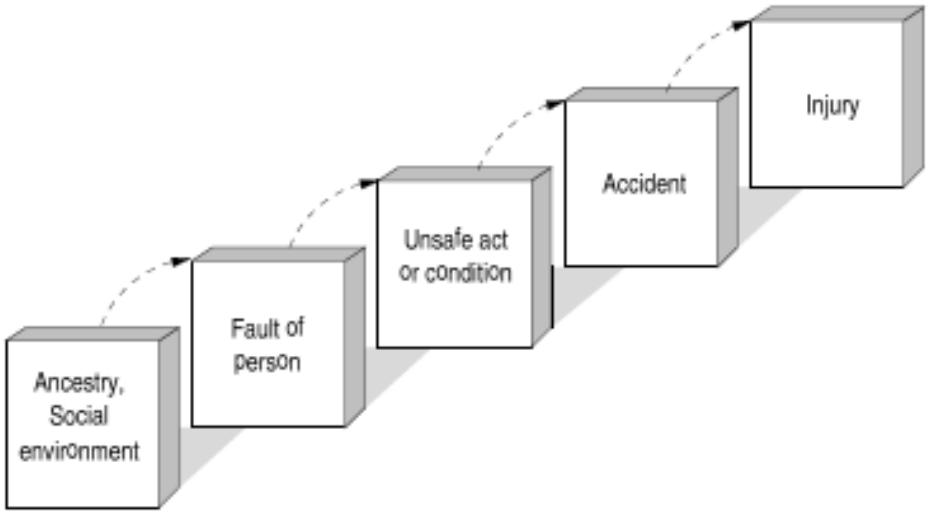

Figure 1. Heinrich's Domino Model of Accidents

In 1976, Bird and Loftus extended the original Domino Model to include the following dominos:

1. Lack of control by management, permitting
2. Basic causes (personal and job factors) that lead to
3. Immediate causes (substandard practices/conditions/errors), which are the proximate cause of
4. An accident or incident, which results in
5. A loss.

In the same year, Adams suggested the dominos should be:

1. Management structure (objectives, organization, and operations)
2. Operational errors (management or supervisor behavior)
3. Tactical errors (caused by employee behavior and work conditions)
4. Accident or incident
5. Injury or damage to persons or property.

The model used by Company X, Reason's Swiss Cheese model [Reason, 1990], in Figure 2 is simply another version of Heinrich's Domino Model with different terminology (slices of cheese rather than dominos) and the idea of barriers or defenses from the process industry. Note that independence of the barriers is assumed and some randomness in whether the "holes" line up. Also, as in the Domino Model, a human (operator) "unsafe act" is seen as the final event while ignoring other types of "active failures" at the time of the accident. Finally, the underlying causes of the events are ignored other than a lack of barriers to prevent their propagation.

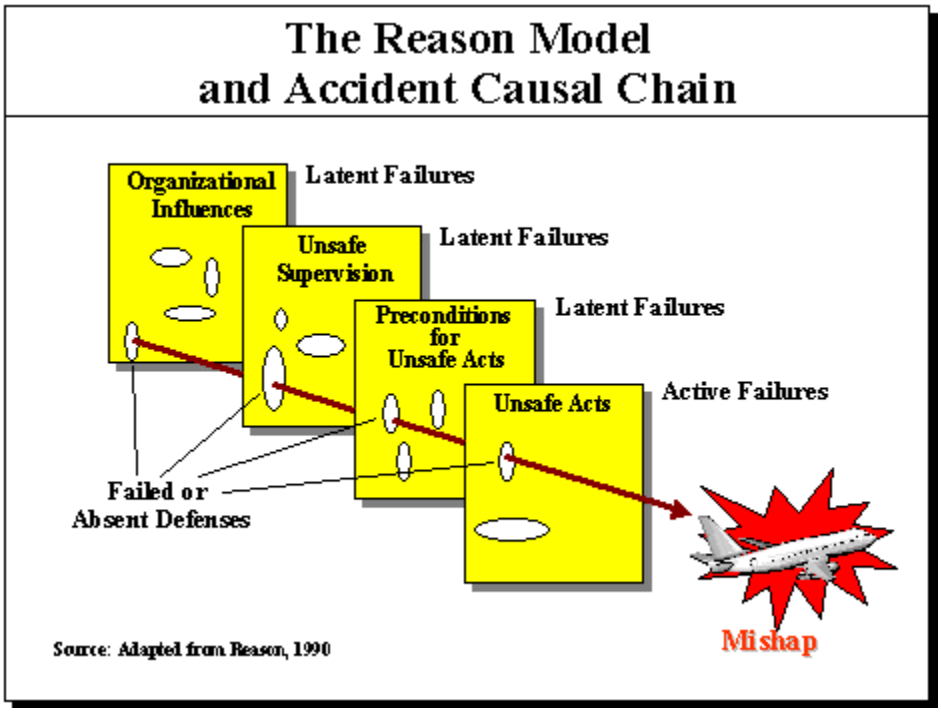

Figure 2: Reason's Swiss Cheese Model

All of these models suffer from oversimplifying accidents as chains of failure events (or dominos or holes in cheese) and not considering the entire accident process. Accidents are often viewed as some unfortunate coincidence of factors that come together at one particular point in time and lead to the loss. This belief arises from a narrow view of the causal time line that focuses only on proximate events or those in a direct chain preceding the event. Systems are not static. Rather than accidents being a chance occurrence of multiple independent events, they tend to involve a migration to a state of increasing risk over time. A point is reached where an accident is inevitable unless the high risk is detected and reduced. The particular events involved are somewhat irrelevant: if those events had not occurred, something else would have led to the loss. This concept is reflected in the common observation that a loss was "an accident waiting to happen."

2.1 Migration Toward States of High Risk

The Columbia Space Shuttle loss provides an example of an accident waiting to happen because the system had migrated to a state of high risk. The proximate cause was foam coming loose from the external tank and damaging the re-entry heat control structure. But many potential problems that could have caused the loss of the Shuttle, such as flow liner cracks, had preceded this event and an accident was avoided in those cases by luck or by unusual circumstances. The economic and political pressures had led the Shuttle program to drift to a high risk state where any slight deviation could have led to the loss [Leveson 2007].

System dynamics can be used to show this. The model below is a simplified model of some of the causal factors in the Columbia loss.

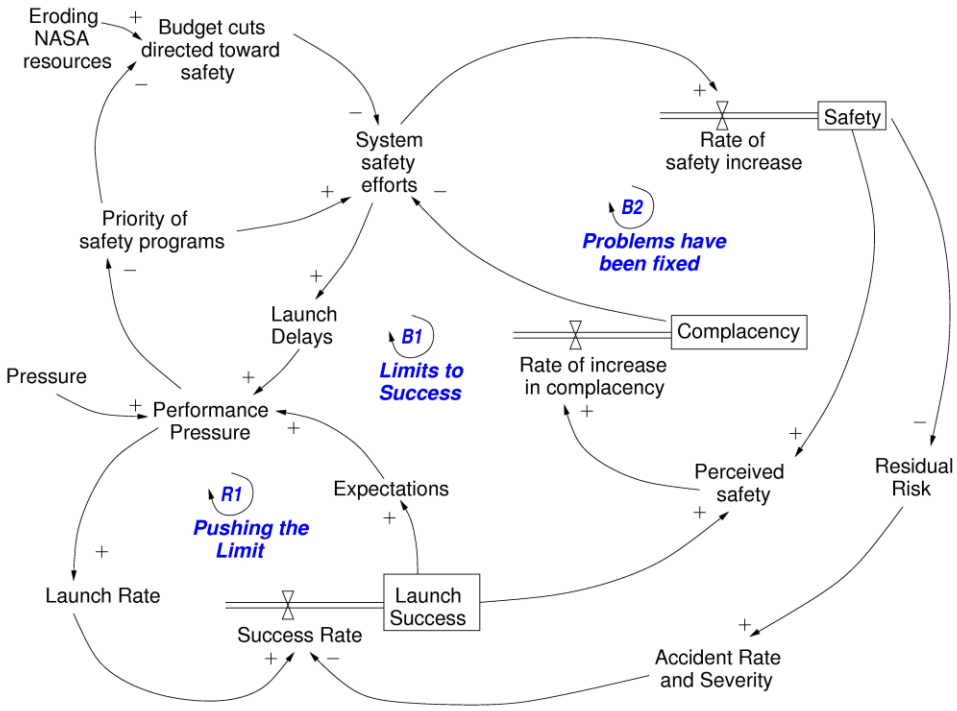

The control loop in the lower left corner labeled R1 or *Pushing the Limit*, shows how as external pressures increased, performance pressure increased, which led to increased launch rates and thus success in meeting the launch rate expectations, which in turn led to increased expectations and increasing performance pressures. This, of course, is an unstable system and cannot be maintained indefinitely—note the larger control loop, B1, in which this loop is embedded, is labeled *Limits to Success*. The upper left loop represents part of the safety program loop. The external influences of budget cuts and increasing performance pressures that reduced the priority of safety procedures led to a decrease in system safety efforts. The combination of this decrease along with loop B2 in which fixing problems increased complacency, which also contributed to reduction of system safety efforts, eventually led to a situation of (unrecognized) high risk. There is one other important factor shown in the model: increasing system safety efforts led to launch delays, another reason for reduction in priority of the safety efforts in the face of increasing launch pressures.

The same factors in the Challenger loss, as noted in the official report on the Columbia accident, contributed to that second Space Shuttle loss because the systemic causes of the Challenger loss had never been fixed. Some of the Challenger proximate events were addressed and behavior changed temporarily, but the systemic causes were still there and the space shuttle program once again migrated to the same high risk level [Leveson 2007].

Understanding and preventing or detecting system migration to states of high risk requires that our accident models focus not on proximate events and human actions but on the entire accident *process*. Processes control a sequence of events and describe system and human behavior as it changes and adapts over time. The overall accident process and system within which the process executes is what causes and explains the proximate events and human actions.

Looking only at the proximate events themselves will lead to identifying and fixing only the symptoms of the underlying problems while allowing accidents from the same systemic causes to recur. Event chain models exclude the systemic factors in accidents that have only indirect or

non-linear relationships to the events. For example, competitive pressures in the industry may have an indirect but important effect on the behavior of operators and managers.

2.2 Systemic Factors in Accidents

Event chain models do not apply to accidents where the cause(s) lies in the interaction among system components, none of which may have failed but each of which may have satisfied its requirements or (for operators) followed the specified procedures. Dominos, event chains, and holes in cheese assume a direct relationship between the events and acts. But causality usually involves indirect, non-linear, and feedback relationships.

Consider Bhopal as an example.[3] The release of methyl isocyanate (MIC) from the Union Carbide chemical plant in Bhopal, India, in December 1984 has been called the worst industrial accident in history: Conservative estimates point to 2000 fatalities, 10,000 permanent disabilities, and 200,000 injuries. The Indian government blamed the accident on human error—the improper cleaning of a pipe at the plant. A relatively new worker was assigned to wash out some pipes and filters, which were clogged. MIC produces large amounts of heat when in contact with water, and the worker properly closed the valves to isolate the MIC tanks from the pipes and filters being washed. Nobody, however, inserted a required safety disk (called a *slip blind*) to back up the valves in case they leaked.

A chain of events describing the accident mechanism for Bhopal might include:

E1: Worker washes pipes without inserting a slip blind.
E2: Water leaks into MIC tank
E3: Explosion occurs
E4: Relief valve opens
E5: MIC vented into air
E6: Wind carries MIC into populated area around plant.

Both Union Carbide and the Indian government blamed the worker washing the pipes for the accident.[4] A different operator error might be identified as the root cause (initiating event) if the chain is followed back farther. The worker who had been assigned the task of washing the pipes reportedly knew that the valves leaked, but he did not check to see whether the pipe was properly isolated because, he said, it was not his job to do so. Inserting the safety disks was the job of the maintenance department, but the maintenance sheet contained no instruction to insert this disk. The pipe-washing operation should have been supervised by the second shift supervisor, but that position had been eliminated in a cost-cutting effort. So the root cause might instead have been assigned to the person responsible for inserting the slip blind.

But the selection of which operator action to label as the root cause (and operator actions are almost always selected as root causes) is not the real problem here—it is the limitations implicit in using a chain of events to understand why this accident occurred. Given the design and operating conditions of the plant, an accident was waiting to happen:

---

[3] For a more complete description of this accident and the limitations of event chain models, see Leveson's new book draft at http://sunnyday.mit.edu/book2.pdf . References are omitted in the description provided in this research report, but many references were used including [Ayres and Rohatgi, 1987; Bogard, 1989; Chisti, 1986; Ladd, 1987].

[4] Union Carbide lawyers argued that the introduction of water into the MIC tank was an act of sabotage rather than a maintenance worker mistake. While this interpretation of the initiating event has important implications for legal liability, it makes no difference in the argument being made here regarding the limitations of event-chain models of accidents.

> "However [water] got in, it would not have caused the severe explosion had the refrigeration unit not been disconnected and drained of Freon, or had the gauges been properly working and monitored, or had various steps been taken at the first smell of MIC instead of being put off until after the tea break, or had the scrubber been in service, or had the water sprays been designed to go high enough to douse the emissions, or had the flare tower been working and been of sufficient capacity to handle a large excursion" [Perrow, 1986, p. 349].

The Bhopal operating manual specified that the refrigeration unit must be operating whenever MIC was in the system: The chemical has to be maintained at a temperature no higher than 5° C. to avoid uncontrolled reactions. A high temperature alarm was to sound if the MIC reached 11° C. The refrigeration unit was turned off, however, to save money and the MIC was usually stored at nearly 20° C. The plant management adjusted the threshold of the alarm, accordingly, from 11° to 20° and logging of tank temperatures was halted, thus eliminating the possibility of an early warning of rising temperatures.

At the Bhopal facility, there were few alarms or interlock devices in critical locations that might have warned operators of abnormal conditions—a system design deficiency. Other protection devices at the plant had inadequate design thresholds. The vent scrubber, had it worked, was designed to neutralize only small quantities of gas at fairly low pressures and temperatures: The pressure of the escaping gas during the accident exceeded the scrubber's design by nearly two and a half times, and the temperature of the escaping gas was at least 80° C. more than the scrubber could handle. Similarly the flare tower (which was supposed to burn off released vapor) was totally inadequate to deal with the estimated 40 tons of MIC that escaped during the accident. In addition, the MIC was vented from the vent stack 108 feet about the ground, well above the height of the water curtain intended to knock down the gas. The water curtain reached only 40 to 50 feet above the ground. The water jets could reach as high as 115 feet, but only if operated individually.

Leaks were routine occurrences and the reasons for them were seldom investigated: Problems were either fixed without further examination or were ignored. A safety audit two years earlier by a team from Union Carbide had noted many safety problems at the plant, including several involved in the accident, such as filter-cleaning operations without using slip blinds, leaking valves, the possibility of contaminating the tank with material from the vent gas scrubber, and bad pressure gauges. The safety auditors had recommended increasing the capability of the water curtain and had pointed out that the alarm at the flare tower from which the MIC leaked was nonoperational and thus any leak could go unnoticed for a long time. None of the recommended changes were made. There is debate about whether the audit information was fully shared with the Union Carbide India subsidiary and about who was responsible for making sure changes were made. In any event, there was no follow-up to make sure that the problems had been corrected.

A year before the accident, the chemical engineer managing the MIC plant resigned because he disapproved of falling safety standards and still no changes were made. He was replaced by an electrical engineer.

Measures for dealing with a chemical release once it occurred were no better. Alarms at the plant sounded so often (the siren went off 20 to 30 times a week for various purposes) that an actual alert could not be distinguished from routine events or practice alerts. Ironically, the warning siren was not turned on until two hours after the MIC leak was detected and after almost all the injuries had occurred and then was turned off after only five minutes—which was company policy.

Moreover, the numerous practice alerts did not seem to be effective in preparing for an emergency: When the danger during the release became known, many employees ran from the

contaminated areas of the plant, totally ignoring the buses that were sitting idle ready to evacuate workers and nearby residents. Plant workers had only a bare minimum of emergency equipment—a shortage of oxygen masks, for example, was discovered after the accident started—and they had almost no knowledge or training about how to handle non-routine events. The police were not notified when the chemical release began; in fact, when called by police and reporters, plant spokesmen first denied the accident and then claimed that MIC was not dangerous. Nor was the surrounding community warned of the dangers, before or during the release, or informed of the simple precautions that could have saved them from lethal exposure, such as putting a wet cloth over their face and closing their eyes. If the community had been alerted and provided with this simple information, many (if not most) lives would have been saved and injuries prevented.

Some of the reasons why the poor conditions in the plant were allowed to persist are financial. Demand for MIC had dropped sharply after 1981, leading to reductions in production and pressure on the company to cut costs. The plant was operating at less than half capacity when the accident occurred. Union Carbide put pressure on the Indian management to reduce losses, but gave no specific details on how to achieve the reductions. In response, the maintenance and operating personnel were cut in half. Maintenance procedures were severely cut back and the shift relieving system was suspended—if no replacement showed up at the end of the shift, the following shift went unmanned. The person responsible for inserting the slip blind in the pipe had not shown up for his shift. Top management justified the cuts as merely reducing avoidable and wasteful expenditures without affecting overall safety.

As the plant lost money, many of the skilled workers left for more secure jobs. They either were not replaced or were replaced by unskilled workers. When the plant was first built, operators and technicians had the equivalent of two years of college education in chemistry or chemical engineering. In addition, Union Carbide provided them with six months training. When the plant began to lose money, educational standards and staffing levels were reportedly reduced.

In the past, UC flew plant personnel to West Virginia for intensive training and had teams of U.S. engineers make regular on-site safety inspections. But by 1982, financial pressures led UC to give up direct supervision of safety at the plant, even though it retained general financial and technical control. No American advisors were resident at Bhopal after 1982.

The financial losses were followed by management and labor problems. Morale at the plant was low. There was widespread belief among employees that the management had taken drastic and imprudent measures to cut costs and that attention to details that ensure safe operation were absent.

These are only a few of the factors involved in this catastrophe, which also include other technical and human errors within the plant, design errors, management negligence, regulatory deficiencies on the part of the U.S. and Indian governments, and general agricultural and technology transfer policies related to the reason they were making such a dangerous chemical in India in the first place. Any one of these perspectives or ``causes'' is inadequate by itself to understand the accident and to prevent future ones. In particular, identifying the root cause of this accident as improper pipe cleaning or sabotage by a low-level worker ignores most of the opportunities for the prevention of similar accidents in the future. Many of the systemic causal factors are only indirectly related to the proximate events and conditions preceding the loss.

When all the factors noted above, including indirect and systemic ones, are considered, it becomes clear that the maintenance worker was, in fact, only a minor and somewhat irrelevant player in the loss. Instead, degradation in the safety margin occurred over time and without any particular single decision to do so but simply as a series of decisions that moved the plant slowly toward a situation where any slight error would lead to a major accident. Given the overall state

of the Bhopal Union Carbide plant and its operation, if the action of inserting the slip disk had not been left out of the pipe washing operation that December day in 1984, something else would have triggered an accident. In fact, a similar leak had occurred the year before, but did not have the same catastrophic consequences and the true root causes of that incident were neither identified nor fixed.

To label one event (such as a maintenance worker leaving out the slip disk) or even several events as the root cause or the start of an event chain leading to the Bhopal accident is misleading at best. Rasmussen writes:

> The stage for an accidental course of events very likely is prepared through time by the normal efforts of many actors in their respective daily work context, responding to the standing request to be more productive and less costly. Ultimately, a quite normal variation in somebody's behavior can then release an accident. Had this `root cause' been avoided by some additional safety measure, the accident would very likely be released by another cause at another point in time. In other words, an explanation of the accident in terms of events, acts, and errors is not very useful for design of improved systems. [Rasmussen, 1997].

In general, event-based models are poor at representing systemic accident factors such as structural deficiencies in the organization, management deficiencies, and flaws in the safety culture of the company or industry. An accident model should encourage a broad view of accident mechanisms that expands the investigation beyond the proximate events: A narrow focus on technological components, operators, and pure engineering activities may lead to ignoring some of the most important factors in terms of preventing future accidents. The accident model used to explain why the accident occurred should not only encourage the inclusion of all the causal factors but should provide guidance in identifying these factors.

2.3 Changing Views of Human Error

In the traditional view of accident causation, human error, particularly operator error, is seen to be the primary cause of incidents and accidents. The solution usually proposed is to do something about humans—fire them, retrain them, discipline them. Alternatively, something can be done about humans in general—marginalize operators by putting in more automation or rigidify their work by creating more rules and procedures [Dekker 2007]. This approach has never been very successful. Operators keep making the same mistakes over and over.

All human actions are influenced by the environment in which they take place. Changing that environment will be much more effective in changing behavior than reward and punishment. Without changing the environment, human error cannot be reduced for long. We design systems in which human error is inevitable and then blame the human and not the system design. Behaviorism, upon which much of occupational safety is based, is an outdated concept in psychology, but it persists in the occupational safety world, and it has influenced process safety in many industries (including the petrochemical industry). Commercial aircraft cockpit design, in contrast, is an example of engineers learning from accidents and making changes to the system design in order to reduce human error.

Part of the problem lies in the event chain approach to accident investigation where it is usually difficult to find an “event” preceding and causal to the operator behavior. If the problem is the system design, there is no proximal event to explain the error, only a decision made during system design. Even if a technical failure precedes the operator action, the tendency is to put the blame on an inadequate response to the failure by the operator. This flawed reasoning occurs several times in the Company X refinery accident report used later as our example. “Latent

failures" simply move the problem into design and management decisions, but still focus on a "failure" or error on the part of some human.

As argued by Rasmussen [Rasmussen, 1997] and others, doing better accident analyses requires shifting the emphasis in explaining the role of humans in accidents from errors to a focus on the mechanisms and factors that shape human behavior, that is, the context in which human actions take place and decisions are made. Modeling human behavior by decomposing it into decisions and actions and evaluating it in isolation from the context in which the behavior takes place is not an effective way to identify what has to be changed to prevent similar errors in the future.

As an alternative, the systems thinking view of human error is that human error is a symptom, not a cause. All behavior is affected by the context (system) in which it occurs: the social context, the value system in which it takes place, and the dynamic work process. To do something about error, we must look at the system in which people work: the design of the equipment, the usefulness of procedures, the existence of goal conflicts and production pressures, etc.

2.4 Hindsight Bias

Another limitation of most accident investigations and analyses is the extensive use of hindsight. Hindsight often allows us to identify a better decision in retrospect, but detecting and correcting potential errors before they have been made obvious by an accident is far more difficult.

After an accident, it is easy to see where people went wrong, what they should have done or not done, to judge people for missing a piece of information that turned out to be critical, and to see exactly the kind of harm that they should have foreseen or prevented. Before the event, such insight is difficult and, perhaps, impossible.

Dekker (2007) points out that hindsight allows us to:

- Oversimplify causality because we can start from the outcome and reason backwards to presumed or plausible ``causes;"
- Overestimate the likelihood of the outcome—and people's ability to foresee it—because we already know what the outcome is;
- Overrate the role of rule or procedure ``violations." There is always a gap between written guidance and actual practice, but this gap almost never leads to trouble. It only takes on causal significance once we have a bad outcome to look at and reason about;
- Misjudge the prominence or relevance of data presented to people at the time;
- Match outcome with the actions that went before it. If the outcome was bad, then the actions leading up to it must have also been bad—missed opportunities, bad assessments, wrong decisions, and misperceptions.

Dekker (2007) distinguishes between *data availability,* which is what can be shown to have been physically available somewhere in the situation and *data observability,* which is what was observable given the features of the interface and the multiple interleaving tasks, goals, interests, knowledge of the people looking at it. The Company X refinery incident provides an example. Our Company X liaison indicated that it was possible for the Board Operator at the refinery to determine the level of liquid in the T-731 tank by "trending the data on the control board" (calling up the trend data), but the operator did not do this. The data was *available* but not *observable* without taking some special action. The operator had no reason to bring up the

trending data as all observable evidence indicated that there was no flow and the operator had other emergencies to attend to at the time.

Knowing the outcome of a sequence of events can easily bias an investigation toward those data points that after the accident we know were significant and showed the real nature of the situation. Without knowledge of the outcome, however, those data points may appear much less significant. Accident investigators need to put themselves in the situation of the operators at the time, with multiple interleaving and overlapping of tasks, other indications and alarms that need attention, and often ambiguous and conflicting data upon which to make decisions.

Consider an example from the Company X refinery SO_2 incident. The report says: "The available evidence should have been sufficient to give the Board Operator a clear indication that Tank 731 was indeed filling and required immediate attention." The report does not identify what that evidence was. From our reading, it seems that the majority of data at the time pointed to the Tank not filling: the operators thought the control valve was closed (the Board Operator had closed it), the flow meter showed no flow, the Outside Operator had reported that he had visually checked and there was no flow, and both of them believed the bypass valve was closed. In retrospect, knowing that there indeed had to have been flow, the report writers may have been able to find some evidence of a flow (a classic example of hindsight bias) or something the operators might have done to get that evidence, but the overwhelming data available at the time to the Board Operator indicated that the tank was not filling and did not require immediate attention. The situation was further complicated by the other alarms that the operators had to attend to at the same time. Other examples of hindsight bias are provided later in this research report.

Hindsight bias occurs because accident investigators know and start from the outcome and trace back to the assessments and decisions that led up to it. Tracing back through the causal flow, it is easy to see where operators could have done something to avoid the outcome—where, as Dekker says, they could have zigged instead of zagged [Dekker 2009]. The question that needs to be answered in the investigation is not what the operators should have done in order to avoid the outcome we now know given our later understanding of the situation. The question needing to be answered is why the operators did what they did given their understanding of the situation *at the time* (which did not include the outcome) and what changes can be made to help them make better decisions under those same circumstances in the future.

Figure 3, from Dekker (2009), illustrates one effect of hindsight bias. Before the accident, the operators and other actors confront the world in its full complexity, with a large number of choices to be made and potential risks to manage. After the accident, that complexity can be boiled down to a single binary decision to see or not to see a particular piece of data, which of course significantly oversimplifies the situation the operators faced at the time.

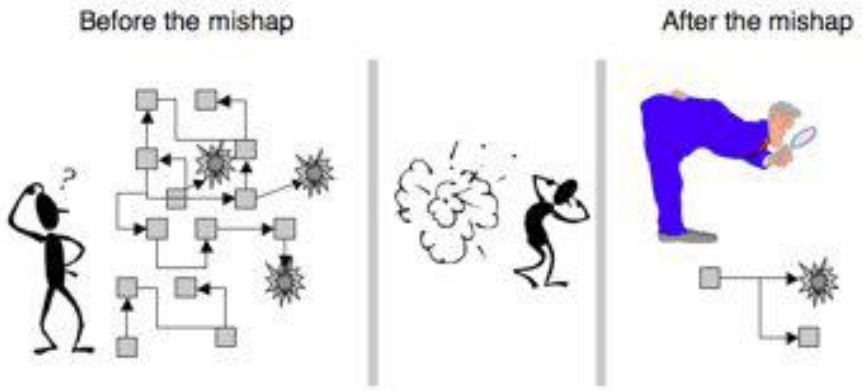

Figure 3. One effect of hindsight bias (Dekker 2009)

Avoiding hindsight bias requires changing our emphasis in analyzing the role of humans in accidents from what they did wrong to why it made sense for them to act the way they did. Examples appear in the refinery SO_2 release report analyzed in depth in Section 3.3

3. Applying Systems Thinking and Systems Theory to Accident Causation

While the traditional approach to process safety critiqued in the previous section comes from industrial safety in the early 1900s, the systems movement in safety engineering is more recent and stems from the perceived inappropriateness of the traditional approaches for the increasingly complex systems we are attempting to build and operate. Leaders of this movement come from the systems engineering and human factors communities and include Rasmussen (system engineering) and Dekker (human factors) in Europe and Leveson (system engineering) and Woods (human factors) in the U.S.

The ideas for this new movement actually predate the academics and researchers and were first applied in military systems after World War II. The potential destructiveness of nuclear and other weapons were increasing so dramatically at that time that focus had to be placed on preventing the first accident. The systems approach to safety has been very effective in this community.

The systems approach to safety engineering is based on systems theory and systems thinking. Systems theory dates from the thirties and forties and arose as a response to the limitations of the classic analysis techniques in coping with the increasingly complex systems starting to be built then [Checkland 1981]. The traditional scientific method breaks systems into distinct parts so that the parts can be examined separately: physical aspects are decomposed into separate physical components while behavior is decomposed into events over time. Such decomposition assumes that such separation is feasible: that is, each component or subsystem operates independently and analysis results are not distorted when these components are considered separately. This assumption, in turn, implies that the components or events are not subject to feedback loops and other non-linear interactions. While these assumptions are reasonable for many of the physical regularities of the Universe, they do not hold for the complex systems we are now engineering [Leveson 2009] nor do they hold for human behavior and social systems [Checkland 1981].

Instead of decomposing behavior into events over time, the systems approach focuses on systems taken as a whole. It assumes that some system properties can only be treated adequately in their entirety, taking into account all aspects from the social to the technical. These system properties derive from the relationships among the parts of the system: how the parts interact and fit together. Thus, the systems approach concentrates on the analysis and design of the whole as distinct from the components or the parts and provides a means for studying emergent system properties such as safety.

Using this approach as a foundation, new types of accident analysis (both retrospective and prospective) can be devised that go beyond simply looking at events and can help identify the processes and systemic factors behind the losses. This information can be used to design controls that prevent hazardous states and detect when risk is increasing before a loss occurs. One such new approach (and the one used in this research) is based on a new, extended causality model called STAMP [Leveson, 2003; Leveson, 2009]

3.1 Introduction to STAMP

STAMP (Systems Theoretic Accident Model and Processes) is a new model of accident causation that uses systems thinking and extends the types of accidents and accident causes considered by traditional accident models.

Using systems theory concepts, safety is an emergent property that results from the enforcement (through system design and operation) of safety-related constraints on the behavior of the system components. An example of a *safety constraint* that was violated at the refinery, for example, is that the acid evolved sulfur dioxide must never be released to the atmosphere for humans to breathe.

Accidents or losses result from unsafe interactions among humans, physical devices, and the organizational and social environment that result in lack of enforcement of the safety constraints. While there are physical controls related to this constraint, there are also organizational and even cultural, societal and governmental controls that assist in enforcing the constraint.

Accidents or unacceptable losses can result not only from system component failures but also from interactions among system components—both physical and social—that violate system safety constraints. Safety is reformulated as a system *control* problem rather than a component *reliability* problem: accidents or losses occur when component failures, external disturbances, and/or dysfunctional interactions among system components are not handled adequately or controlled—where controls may be managerial, organizational, physical, operational, or manufacturing—such that required safety constraints on behavior are violated. The actual process that led to the lack of control (or accident) can be very complex and may involve indirect, non-linear, and feedback relationships among the events and the system components.

Every system includes a set of not only physical controls to ensure safety but also social, organizational, and managerial controls. To understand why an accident occurred, it is necessary not only to determine why the *physical* controls were inadequate, but why the *social, organizational, and managerial controls* (in development and operations) were not able to prevent the loss event. These controls are embodied in the *hierarchical safety control structure*. Figure 4 shows an example of a hierarchical safety control structure for a typical regulated industry, such as commercial aircraft, in the U.S.

Feedback control loops exist between many of the safety control structure components to provide the information necessary for the component to effect adequate control (Figures 4 and 5). For example, operations management provides control actions such as work instructions and procedures and receives feedback in the form of change requests, audit and inspection reports, problem reports, incidents, etc. Higher level controllers may provide overall safety policy, standards, and procedures and get feedback about their effect in various types of reports, including incident and accident reports. The feedback provides the ability to learn and to improve the effectiveness of the safety controls.

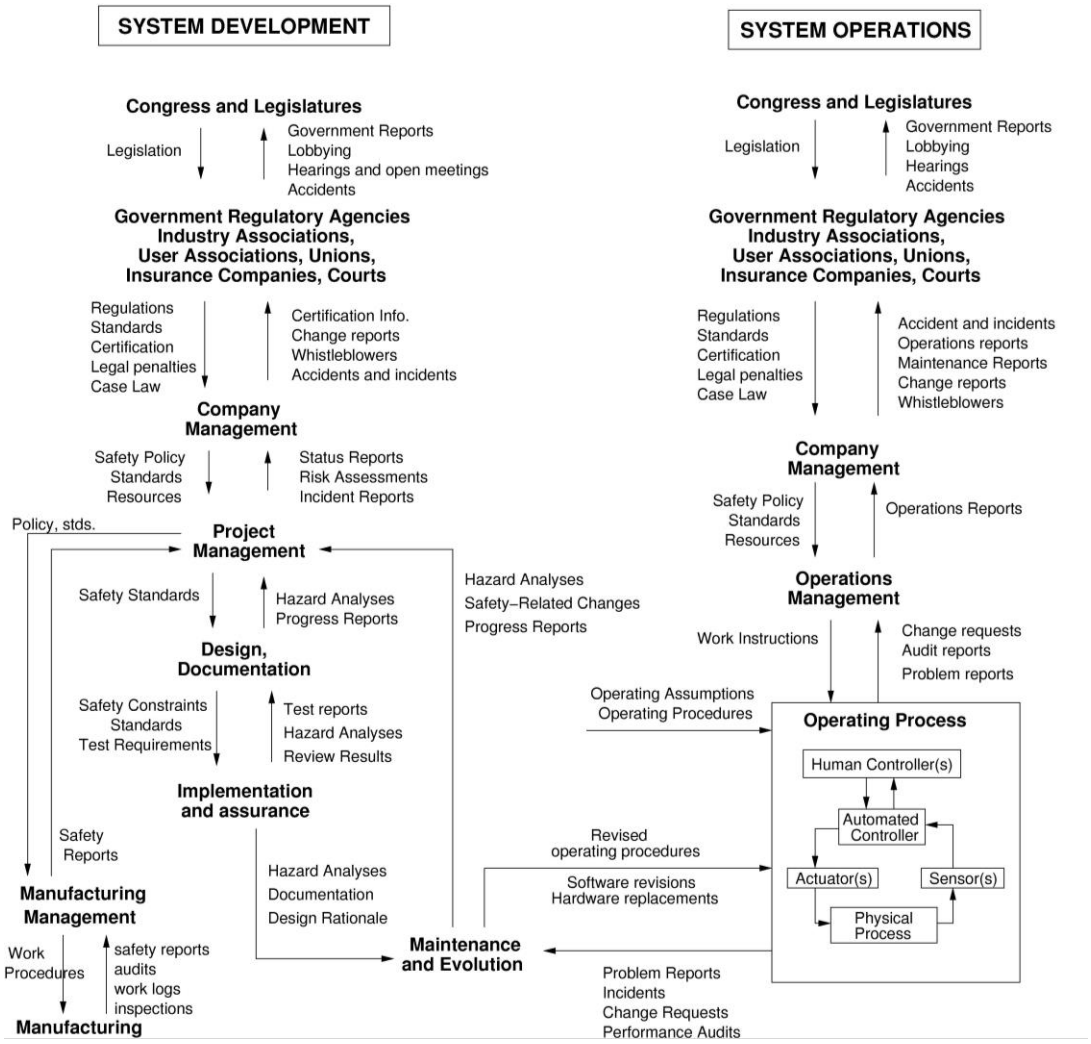

Figure 4. An Example Safety Control Structure

One other important concept is needed. In basic system (and control) theory, in order to provide effective control, the controller must have an accurate model of the system it is controlling. For human controllers, this is usually called the m*ental model.* This process model or mental model is used to determine what control actions are necessary to keep the system operating effectively.

The process model includes assumptions about how the controlled process operates and the current state of the controlled process. Accidents in complex systems often result from inconsistencies between the model of the process used by the controller and the actual process state. The flawed process model leads the controller to provide inadequate control. Usually, these models of the controlled system become incorrect due to missing or inadequate feedback and communication channels. In the Company X overflow accident, for example, the Board Controller was missing information about the level of SO_2 in the tank due to a non-functioning high-level alarm and flow meter. Had he had this information, his behavior might have been very different.

.

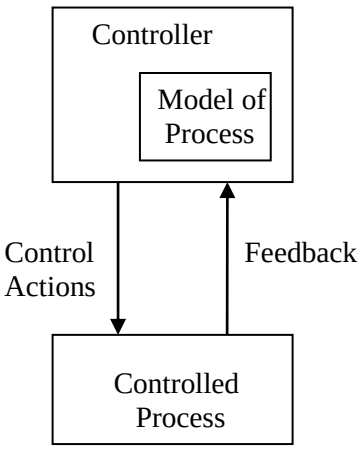

Figure 5. Every controller contains a model of the controlled process.

Part of understanding causality in accidents and why people did the things they did requires understanding their flawed mental models (the information they had at the time), as well as the context in which the decisions take place.

When losses occur, either some or all of the hierarchical safety control system components did not fulfill their responsibilities or the overall design of the system (the design of the components and their responsibilities) did not adequately fulfill the system goals and needs to be redesigned (or both). Analyzing accident causality then involves identifying the components of the safety control structure that did not exercise adequate control and, more important, why they were unable to fulfill their responsibilities.

The hierarchical safety control structure is usually not static. Major accidents are often the result of a dynamic process that started before, sometimes years before, the actual physical loss events [4]. Refer to the Columbia accident system dynamics model on Page 3. Organizations and sometimes even whole industries move slowly toward a state of high risk behavior and decision-making under extreme performance and competitive pressures that makes an accident almost inevitable unless detected and corrected in time. Nobody intends to harm other people, but economic stresses may lead to taking larger risks or inadequately executing responsibilities. In STAMP, that dynamic process is conceived as resulting from an adaptive feedback function that fails to maintain safety as performance changes over time to meet a complex and changing set of goals and values.

It is the responsibility of the hierarchical safety control system to prevent migration to unacceptably high states of risk (i.e., unacceptable system component behavior) or to detect when it is occurring and respond appropriately. Understanding why accidents occur in complex socio-technical systems requires understanding the context in which decision making takes place, particularly those factors that militate against a controller providing the control necessary to successfully fulfill its responsibilities. For example, production pressures without corresponding system design features to assist in making better decisions in the face of such pressures (e.g., better ways to assess the risk involved in the actions taken or procedures to follow to reduce that risk) can play an important role in the flawed decision making leading to an accident.

To summarize, using STAMP as the basis for accident investigation and analysis involves looking not only at the events that occurred prior to a loss in order to determine why it occurred and how to prevent future occurrences, but also examining the entire dynamic accident or loss process, i.e., why the overall safety control structure did not enforce constraints on the behavior

of the system components that would have prevented the loss. It may have been inadequate from the beginning or it may have been effective at one time but grew less effective over time.

STAMP assumes that accidents occur due to violation of safety constraints. These violations may result from environmental disturbances or conditions, system component failures, or unsafe interactions among the system components. Inadequate safety control actions can be traced to:

1. A lack of designed controls, at the physical, social, and organizational levels
2. Inadequate operation of the existing controls, perhaps due to:
   a. Social and political factors (context)
   b. Controller process models that do not match the state of the process being controlled and lead to inadequate control actions on the part of the controller.
3. Degradation and changes in the safety-control structure over time, both planned and unplanned.
4. Inadequate coordination of safety-control actions among multiple controllers.

A nice simple diagram to illustrate this model is difficult, but the figure below is my attempt at it.

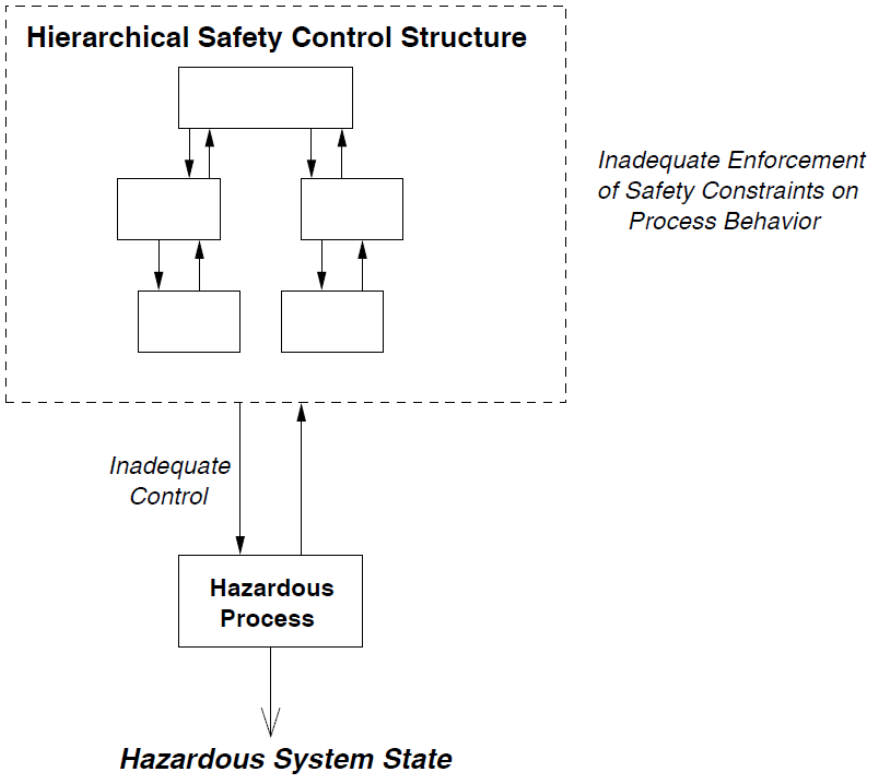

Figure 5b: An accident occurs when the safety control structure does not enforce the safety constraints on the behavior of the controlled process, leading to a hazardous system state.

To completely understand the cause of accidents and to prevent future ones, the system's hierarchical safety control structure must be examined to determine why the controls at each level were inadequate to maintain the constraints on safe behavior at the level below and why the events occurred. A complete description of STAMP and how to use it to analyze accidents is beyond the scope of this report—it can be found in [Leveson 2012]. The analysis of the refinery Tank 731 overflow incident using STAMP in Section 3.3 should be understandable without it.

3.2 The Role of Blame in Accident Analysis

There are usually two reasons for conducting an accident investigation: (1) to assign blame for the loss or (2) to understand why it happened so that future accidents can be prevented. When the goal is to assign blame, the backward chain of events considered often stops when someone or something appropriate to blame is found, such as the maintenance worker at Bhopal or the operators and failed gauges and indicators at the refinery. But the selected initiating event in the proximal event chain usually provides too superficial an explanation of why the accident occurred to prevent similar incidents in the future. In addition, for a variety of reasons (beyond the scope of this report), blame is usually assigned to operators and not to management or organizational and cultural flaws. A final limitation of focusing on blame as a goal of the accident investigation process is that the investigation process is hindered as those on whom blame may be placed are unlikely to be forthcoming and honest about their possible role in the loss events.

Blame is not an engineering concept; it is a legal and moral one. Usually there is no objective criterion for distinguishing one factor or several factors from the many other factors contributing to an accident. When learning how to engineer safer systems is the goal rather than identifying who to punish and establishing liability, then the emphasis in accident analysis needs to shift from *cause* (in terms of events or errors), which has a limiting, blame orientation, to understanding accidents in terms of *reasons*, i.e., why the events and errors occurred. The majority of accident reports in many industries stop after assigning blame—usually to the operators—and never get to the root of why the accident occurred, e.g., why the operators made the errors they did and how to prevent such errors in the future. Accident reports also rarely identify the systemic factors that contributed to the loss.

When trying to understand human contributions to accidents (whether the humans are the operators, managers, engineers, etc.), just as with overcoming hindsight bias, it is more helpful in learning how to prevent future accidents by focusing *not* on what the human did "wrong" but on why it made sense for them to behave that way under those conditions. Most people are not malicious, but simply trying to do the best they can under the circumstances and with the information they have. Understanding why those efforts were not enough will help in changing features of the system and environment so that sincere efforts are more successful in the future. Focusing on assigning blame contributes nothing toward achieving this goal and may impede it.

Because our goal is to learn from events in order to prevent future incidents and accidents, the emphasis in our accident analysis approach is on understanding why the accident occurred and not on assigning blame. That implies identifying *all* the factors involved and understanding the relationship among these causal factors. That information (feedback) can then be used to generate recommendations for preventing losses in the future. Building safer systems will be more effective when we consider all causal factors, both direct and indirect.

There also is no way to determine whether some factors are more important than others. In the STAMP view of accidents as a process, there is no such concept. All factors contributing to the process were necessary for the loss to occur and, more important, we can learn from all of them in terms of redesigning the safety control structure to be more effective in the future.

3.3 Company X Refinery Accident/Incident Analysis using STAMP

To determine whether these new accident analysis ideas would apply to Company X and the process industry, we analyzed a recent incident at a Company X refinery. In this section we present our analysis. In Section 3.4, the results and final recommendations are compared with the accident report generated using Company X's standard accident analysis methods, which are similar to those used throughout the process industry and other industries.

STAMP-Based Analysis of the SBS Tank 731 Overflow

Description of the Physical Process:

In the SBS unit, tail gas is burned at 1200° F with excess air and natural gas in the tail gas combustor F-700. This converts the H_2S to SO_2. [A physical diagram could be inserted here, but we did not have one that we could use.]

Hot gas effluent from F-700 is cooled in the waste heat boiler E-701. Effluent gas from the boiler enters the venturi quench tower V-703 where it is quickly cooled by direct contact with a 30% sulfuric acid solution to approximately 181° F. The 30% acid solution comes from the bottom of the T-704 quench separator and is pumped to V-703 via the quench circulating pumps P-704A/B/C in a continuous loop. The acid concentration is maintained at 30% by purging a small (< 2 gpm) slip stream to Tank-731 acid tank and replacing it with water to maintain the process temperature near 181° F. The acid from Tank 731 is then removed either by draining to the process sewers (used as Lakefront pH control) or used as a product elsewhere. As the acid enters Tank 731, it is saturated with SO_2. The design of Tank 731 allows for the gas that is entrained in the acid to degas off to another tower, T 707, where it can be further treated. On Thursday August 7th, when tank 731 tank overfilled, there was not sufficient residence time for the SO_2 entrained in the solution to degas properly. As a result, the acid evolved sulfur dioxide to atmosphere as it overfilled to the ground.

Events:

The analysis starts, like any accident analysis, with identifying the proximate events including the physical failures and operator actions (or missing actions) related to the loss. But stopping after identifying these, often the end point in accident investigation, usually leads to attributing the cause to operator error, which, as stated earlier, does not provide enough information to prevent accidents in the future. The operators may be fired or reprimanded, subjected to additional training, or told not to make the same mistake in the future, none of which lead to long term prevention of the same behavior if problems exist in the other parts of the safety control system design. It also leads to identifying and fixing specific hardware design flaws, e.g., the redesign of a relief valve or the replacement of a flow meter, but not the flaws in the engineering design and analysis process or the maintenance issues that led to that particular manifestation of a flawed design. Examining the rest of the control structure will provide more information about the flaws in the larger company safety management structure that need to be fixed.

The events below are from the Company X accident report (see Appendix).

08:33 - Board Operator attempts to open control valve F-47706, to begin an acid drawdown from the quench recirculation system. The flow meter does not indicate a flow, so the Board Operator calls the Outside Operator to check and see if the manual block valves at the control valve station are closed. Note: The Board Operator stated/believed that the block valves are normally left in the open position to facilitate conducting this operation remotely.

09:11 - Outside Operator finds that the manual block valves are lined up (open) and has the Board Operator open the control valve to different settings in an effort to troubleshoot the situation. Outside Operator also sees no indication of flow on the flow meter and makes an

effort to visually verify that there is no flow. He then begins to work (open and close) the manual block valves in an effort to fix the problem. Note: Process control data indicates that the tank level indicator begins to show an increase in the tank liquid level around this time. Acid level in the tank is approximately 7.2 ft. at this time. Per the interviews, neither the Board Operator nor the Outside Operator had any discussion about opening the bypass valve. The Board Operator did not call for the bypass valve to be opened, and the Outside Operator states that he did not open the bypass valve.

09:25 - Outside Operator finishes his effort to work (open and close) the manual block valves. He reports having heard a clunking sound and thought something might have "broke loose" so he asks the Board Operator to try opening the control valve again. Outside Operator still sees no flow on the flow meter but does not make another effort to visually verify this condition. Note: the tank level is now at approximately 7.7 ft. at this time. Outside Operator gets a call to perform other unit duties and tells the Board Operator to call him when he wants to try again. Outside Operator leaves the manual block valves at the control valve station in the open position. Board Operator leaves the control valve in the closed position (confirmed by process control data).

09:37 - Tank 731 high level alarm sounds in the control room Tank level is at approximately 8.5 ft. Board Operator acknowledges the alarm. About a minute and a half later the alarm is disabled.

09:49 - Alarm B45002L on the Beavon-Stretford unit sounds, indicating an Emergency RGG-One Fire Eye went out. This event and alarm is associated with on-going unit operations to move Pit Sweep from the SBS to Beavon-Stretford.

09:50 - Tank 731 appears to overflow (i.e., chart flat lines).

10:00 - SO_2 alarm (A47710) sounds at 4 ppm, but quickly climb to 25 ppm (maximum instrument reading). At about the same time, emergency alarm B45002LL at the Beavon-Stretford goes off, indicating both fire eyes on the RGG went off, causing it to trip. Board Operator contacts Outside Operator via radio and asks him to check it out at the Beavon-Stretford unit

~10:25 – Based on interviews, at approximately this time exposed workers make their way to area southeast of affected area and report odor and irritation problems to their Job Rep.

~10:31 – Based on interview with Outside Operator, at approximately this time the manual block valves around the control valve were closed by Outside Operator.

10:48 – Company X ambulance requested

10:54 - Unit evacuation alarm sounded by unit asset supervisor

11:33 - SO_2 concentration drops below 4 ppm (alarm set-point) on analyzer (A47710).

13:18 - SO_2 concentration is non-detectable at analyzer (A47710).

Safety Control Structure

The safety control structure (process safety management system) consists of the controls that have been implemented to prevent hazards. In order to understand why the accident (the events) occurred using systems thinking and treating safety as a control problem, it is necessary to determine why the controls created to prevent it were unsuccessful and what changes are necessary to provide more effective control over safety.

Figure 6 shows our best reconstruction of the safety control structure at Company X from what we could glean in the accident report and from our knowledge of the company and the process industry in general. It is very general and we realize there are almost surely many errors in it, but this was the best we could do. It is adequate for our demonstration.

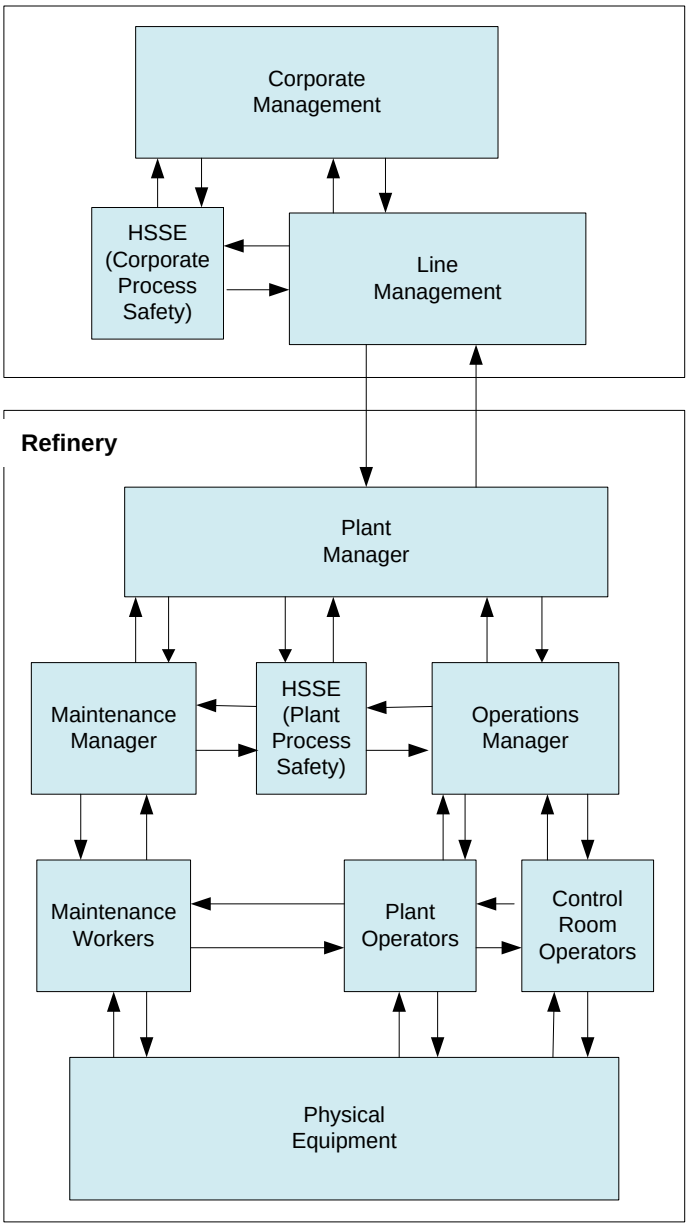

Figure 6. A Simplified Safety Control Structure for Company X Refineries

Each of the components in the safety control structure has particular responsibilities with respect to safety. As we have very little information about the actual responsibilities assigned to each group at Company X, we inferred responsibilities that seemed reasonable to us. For example, we assumed that the maintenance group was responsible for maintaining safety-critical equipment (as opposed to some special group being given this responsibility) and that HSE (safety engineering) was responsible for performing hazard analyses and risk assessments and producing safety-related operator procedures. An accident analysis using STAMP involves determining whether these responsibilities were carried out and, if not, why not. If we have inadvertently mis-assigned responsibilities with respect to Company X and the refinery, it has little impact on the analysis as those responsibilities should be assigned to someone. The goal is not to determine blame but to identify weaknesses in the safety control structure and the changes that need to be made to prevent future losses.

The component responsibilities are identified in the analysis of each component that follows.

Physical Process Analysis

After the control structure has been constructed, the next step is to examine each component, starting with the lowest physical controls and progressing upward to the social and political controls.

The analysis of the physical controls does not differ significantly from that done in most accident analyses. We start by looking at the physical plant safety controls at the bottom of the safety control structure and work upward. We include only those controls related to the specific SO_2 event although we tried to include general responsibilities as far as we could deduce them. The report is not as complete as it would have been if we had collected more information than that provided in the original accident report and in the response to a few questions we were able to ask about the incident. But it serves to demonstrate the power of this approach compared to a standard accident analysis.

SO_2 is a colorless gas with a pungent odor. The acid evolved sulfur dioxide to atmosphere as it overfilled to the ground. SO_2 causes respiratory tract, skin, and eye burns at high concentrations. Workers who were exposed during the incident reported feeling a burning sensation in their eyes, nose, throat, and lungs.

The design of Tank 731 allows for vapor that is entrained in the acid to degas off to another tower, T707, where it can be further treated. As this tank overfilled, SO_2 entrained in the solution did not have sufficient residence time to degas properly. As a result, the acid evolved sulfur dioxide to atmosphere as it overfilled to the ground. The drain was blocked and thus the SO_2 could not flow into the process sewer.

The plant safety equipment (controls) was designed as a series of barriers to protect against runaway reactions, protect against inadvertent release of toxic chemicals or an explosion (uncontrolled energy), convert any released chemicals into a non-hazardous or less hazardous form, provide protection against human or environmental exposure after release, and provide emergency equipment to treat exposed individuals. It appears that the refinery has the standard types of safety equipment installed.

Physical Controls and Safety Related Equipment

Requirements (roles/responsibilities): Provide physical protection against hazards (protection for employees and others within the vicinity);

1. Protect against runaway reactions
2. Protect against inadvertent release of toxic chemicals or explosion
3. Convert released chemicals into a non-hazardous or less hazardous form
4. Contain inadvertently released toxic chemicals
5. Provide feedback to operators and others about the state of safety-critical equipment
6. Provide indicators (alarms) of the existence of hazardous conditions
7. Provide protection against human or environmental exposure after release
8. Provide emergency treatment of exposed individuals

Emergency and Safety Equipment (controls): Only those related to the Tank 731 overflow and subsequent events are included.

- Flow meter and level transmitter
- Block valves, bypass valve
- SO_2 alarm
- High level alarms
- SO2 alarm (analyzer): Strobe light
- Unit evacuation alarm
- Drain from containment area to process sewers
- Process vent routed to T-707 from T-731.
- Overflow pipe with gooseneck
- RV

Failures and Inadequate controls: (the links below refer to the requirements above)

- SO_2 released to atmosphere (→ 2)
- Control flow valve may have stuck open (→ 2)
- Level transmitter L47731A for Tank 731 was not working properly. Readings had been erratic for a year and a half. This meant that one of the high level alarms was effectively disabled. (→ 5)
- Flow meter FT47706 was not working properly (→ 5)
- Drain to emergency containment sewer clogged. (could not send excess gas to safe containment area) (→ 4)
- Alert for harmful release of toxic SO_2 is visual and could not be seen by workers in path of released gas. SO_2 analyzers on the SVS alarm trigger flashing strobe lights on the unit, but no audible alarm so they are only effective if they are within the workers line of sight. Several of exposed workers were over 100 yards from the unit and were not able to see the flashing lights. (Because SO2 is a gas, it has the potential to travel away from the unit and around objects to reach workers who may not be able to see the flashing strobe lights.) (→ 5)

Physical Contextual Factors:

- Wind was from NNE at about 9 mph.

(The links in the Failures and Inadequate Controls section refer to the requirements listed above it.)

While there was a reasonable amount of physical safety controls provided, much of this equipment was inadequate or not operational, as noted on the previous page. For example, the level transmitter on the tank was not working properly, the flow meter was not working properly, and the drain to the emergency containment sewer was clogged. The report does not go into why the sewer was clogged and whether this was a common occurrence. While the sewer would not have prevented the overflow, it seems to us that the pooling of the SO_2 in the containment area contributed to the amount of gas in the air and the level of exposure of the workers to this gas.

The failures and inadequate controls raise many questions not included in the original accident report. Were non-functioning or inadequately functioning critical physical controls unusual or was it common at the plant? What types of operational policy exists (in written form or implicit in the culture) about operating a unit with safety equipment out of order? Is a risk assessment done when operations occur under those conditions? What types of inspections are performed on safety-critical equipment? How is safety-critical equipment identified? What was the maintenance policy and why was safety-critical equipment non-operational or operating erratically for relatively long periods of time? Most of these questions are raised later under each of the appropriate safety control structure components.

Safety Control Structure Analysis

The next step in the accident analysis is to determine why the physical system events occurred by examining the higher level control components. The analysis uses the identified responsibilities to determine any inadequate control provided by the component that might have contributed to the loss event. For each component in the safety control structure, the following are identified:

- The safety-related responsibilities of the component
- The component's control actions related to the loss
- The context in which those actions occurred
- Potential process model flaws that led to the inadequate control actions by the component

We start from the assumption that most people have good intentions and do not purposely cause accidents. We need to understand, then, why they did the wrong thing in the particular situation in which they found themselves. In particular, we are interested in the contextual or systemic factors and flaws in the safety control structure that influenced that behavior. For example, the operator or manager may have a mental model that is inconsistent with the real process state and therefore they provide inadequate control. At this refinery, for example, the Board Operator thought that there was no flow because the flow meter showed no flow. There also was missing feedback about the state of the controlled process such as the level of liquid in the tank. The same is true at higher levels of the control structure. For example, the accident report says that there had been previous instances of the units not being properly evacuated in situations where workers may have been at risk. Did the higher levels of management responsible for safety at the Company X refineries know about these instances and just not do anything about it or were they missing feedback about inadequate evacuation procedures and behavior at the refinery? If they had known, we suspect they might have done something about it before the Tank 731 overflow although we cannot know this for sure. We can determine, however, the changes that need to be made to improve the chances that it will not happen again.

In order to minimize hindsight bias, our emphasis in explaining human actions is on understanding why it made sense at the time for the people to act the way they did. We get this information by looking at the context in which the actions took place and the mental (process) model flaws that contributed to the inadequate control.

Board Operator

***Safety-Related Responsibilities*:**

- Provide control actions that avoid hazardous states in the plant.
- Respond appropriately when a potentially hazardous condition or event arises.
- Report incorrectly functioning equipment.

***Context*:**

Related to Tank Level:

- Flow meter was broken: Indicated no flow
- Level transmitter L47731A and its high-level (7.5ft) alarm were not functioning properly. The level transmitter had been erratic since January 2006 but a work order was not written to repair it until July 2008.
- Level transmitter L47731 and its high-level (8.5ft) alarm were functioning.
- Level transmitters gave conflicting information regarding tank level.

Related to Risk and Procedures:

- No written unit procedure for responding to an SO_2 alarm. The "standard response" to an SO_2 alarm on the SBS unit is to have operator conduct a field assessment of the situation.
- No written procedure for ordering evacuation when an SO_2 alarm sounds or criteria established for the level of SO_2 that should trigger an evacuation alarm.
- Unit training material does contain information on the hazards of SO_2, including IDLH information, but this information has not been instituted in standard operating/emergency procedures.
- Block valves normally left open to facilitate remote operations.

Related to Alarms:

- Distracted by other duties related to transferring the Pit Sweep from the SBS to the Beavon-Stretford, which demanded his attention during the time of this incident.
- An alarm indicating that the RGG fire eye went out sounded just before the operating level sensor reached maximum value (no alarm generated).
- Ten minutes later, an alarm indicating that both RGG fire eyes went out sounded at approximately the same time as the SO_2 alarm. This means that multiple alarms were going off at the same time.
- Previous SO_2 alarms were attributed to minor releases that did not require a unit evacuation. Such alarms occur approximately once a month.
- None of alarms were designated as critical alarms, "which may have elicited a higher degree of attention amongst the competing priorities of the Board Operator."
- Upper detectable limit on SO_2 analyzers is 25 ppm. During the incident, analyzer A47710 maxed out at 25 ppm almost instantly, making it impossible to determine the actual SO_2 concentration during the incident.
- There is no established criteria in a written procedure for what SO_2 levels and/or alarms constitutes an emergency condition that should trigger sounding the evacuation alarm. There is also no specific SO_2 ppm threshold for unit evacuation.
- In past, units were not evacuated by blowing the horn, but rather by operations personnel walking through the unit and stopping work.
- Evacuation alert must be sounded by hand, no written procedure for doing so.

***Control actions related to the loss*:**

Control Valve:

- Opened control valve at 8:33. Flow meter does not indicate a flow.
- At 09:25 opened control valve to 44% of open position for about 2 minutes, then closed it.

Alarm:

- At 9:37-9:39: Board Operator acknowledges and disables tank 731 high level alarm. Takes no other action.
- Did not communicate the high-level tank alarm to the Outside Operator
- Delayed sounding evacuation alarm until 54 minutes after the SO_2 was detected by analyzer A47710.

***Process Model Flaws*:**

Tank:

- Believed there was no flow into tank.
- Did not know that flow meter was not functioning.
- Did not know that level transmitter L47731A was not functioning.
- Did not know that level transmitter L47731 indicated a rising tank level. Believed the liquid level was "tickling" the sensor and triggering a false alarm.

SO_2 Risk

- Was not aware of hazards to human health due to SO_2 exposure.
- Was not aware of risk of SO_2 release via tank overflow.
- Did not know the process sewers were clogged, which would increase the amount of SO_2 in the air in case of a tank overflow.
- Did not know SO_2 release warranted evacuation order.
- Was not aware of actual levels of SO_2 but only maximum reading of 25 ppm.

Plant:

- During the incident, did not appear to have a "full awareness" of all the work being performed by people near the unit.

The operator clearly was confused about the level and flow of SO_2. One way to help clarify why this occurred is to examine the information the operator had at each point in time where critical decisions were made. The Board Operator first attempted to open the control valve to begin the acid drawdown at 8:33. The operator's process model at the time looked like the following:

Board Control Valve Position: *open*	Flow Meter: *shows no flow*
Manual Control Valve Position: *open*	Flow: *none*
Bypass Valve: *closed*	SO_2 alarm: *off*
Level on tank: 7.2 feet	High level alarm: *off*

We have left the two redundant level indicators off the process model because we are not sure what information they actually provided. The report says that one was behaving erratically and providing fluctuating readings at the time of the incident. If at least one did provide accurate

information, then the question that needs to be answered is why that information was not used by the Board Operator. There are so many different possible explanations for this omission that it would not be helpful to guess; the reason needs to be carefully investigated. We left it out of his process model as apparently the Board Operator did not use that information or discounted it in preference to the other information he had. We were told that the operational level indicator did indicate 7.2 feet but we do not know if or how that reading changed over time or whether the Board Operator even looked at it.

At this time the flow meter shows no flow, so the Board Operator calls the Outside Operator to see if the manual block valves are perhaps closed. He believes they are usually left open to facilitate remote operation.

Between 9:11 and 9:25, the Outside Operator works the manual block valves to see if he can fix the problem. He hears a clunk and asks the Board Operator to try opening the control valve again. He tries, but they both still see no flow on the flow meter. The Outside Operator gets a call to do something else and leaves the manual valve in the open position. The Board Operator's process model at this time is:

Board Control Valve Position: *open*	Flow Meter: *shows no flow*
Manual Control Valve Position: *open*	Flow: *none*
Bypass Valve: *closed*	SO_2 alarm: *off*
Level in tank: 7.2 feet	High level alarm: *off*

At 9:25, when the Outside Operator left, the Board Operator closes the Board Control Valve. There is no way for the Board Operator to get confirmation that the valve has actually closed—the valve was not equipped with a valve stem position monitor so the Board Operator only knows that a signal has gone to the valve for it to close. The operators in many accidents, including Three Mile Island, have been confused about the actual position of the valve due to similar designs. While obtaining actual position information adds extra expense, the cost may be justified for critical information.

Board Control Valve Position: closed	Flow Meter: *shows no flow*
Manual Control Valve Position: *open*	Flow: *none*
Bypass Valve: *closed*	SO_2 alarm: *off*
Level in tank: 7.2 feet	High level alarm: *off*

The report writers calculate that the flow actually started sometime between 9:11 and 9:25. So clearly the Board Operator's process model was now incorrect but there was no feedback or information available to the Board Operator to use to detect this mismatch. The report calculates that the actual level in the tank was 7.7 feet at 9:25. There is an alarm at 7.5 feet, but it was not working at the time. In answer to our questions, we were told that the operator did not know that the alarm was not working so it seems reasonable to us that the operator would conclude that the level in the tank had not risen above 7.5 feet given that he believed there was no flow into the tank and the 7.5 foot alarm had not sounded.

At 9:37 (twelve minutes later), the Tank 731 high-level alarm sounds in the control room. The tank level must have been at 8.5 feet at the time as that is the level of the second alarm (which was working). The Board Operator's process model now is:

Board Control Valve Position: closed	Flow Meter: *shows no flow*
Manual Control Valve Position: *open*	Flow: *none*
Bypass Valve: *closed*	SO_2 alarm: *off*
Level in tank: 7.2 feet	High level alarm: on

Note that the information the Board Operator has is that the input valve is closed, there is no flow into the tank, and the bypass valve is closed. He acknowledges the alarm, assumes it is spurious, and one and a half minutes later turns the alarm off. He takes no action because he says he believed the liquid level was "tickling" the sensor and triggering a false alarm. In response to our questions, we were told that it is possible for the Board operator to determine the liquid level by "trending the data on the control board" (calling up the trend level data) but the operator did not do this. From what we can determine, the operator must request the trend data—it is not normally shown on the control panel by default. The Board Operator at this point had no reason to believe he needed it. His process model now at 9:39 is:

Board Control Valve Position: closed	Flow Meter: *shows no flow*
Manual Control Valve Position: *open*	Flow: *none*
Bypass Valve: *closed*	SO_2 alarm: *off*
Level in tank: 7.2 feet	High level alarm: off

At 9:49, an alarm on another unit sounded and the operator switched his attention to that. At 9:50 Tank 731 overflowed. At 10:00, the SO_2 alarm sounded at 4 ppm but quickly climbed to 25 ppm, which was the maximum reading. At the same time, there was another emergency alarm at another unit. The Board operator calls the Outside Operator to tell him to check out the other unit.

At this point we need more information than we were able to obtain from the accident report to try to determine why later actions or non-actions occurred. It is clear, however, that the Board Operator's process model was incorrect. To prevent such a mismatch in the future, changes need to be made in the system in order to get better information to him. Such changes may involve different training, different procedures, different feedback channels, control room changes, engineering changes, maintenance changes, etc. Simply saying the operator made a mistake is not helpful without a careful analysis of why and how to prevent it in the future.

There are several instances of hindsight bias in the accident report relevant to the operators' role in the incident. One example was noted earlier. The report says:

> "The available evidence should have been sufficient to give the Board Operator a clear indication that Tank 731 was indeed filling and required immediate attention."

The report does not expound on what that evidence was. From our reading, the majority of the evidence the operators had did not indicate the tank was filling and, in fact, indicated the opposite as shown above. The operator did not use his data trending tool, but there was little reason to do so before the actual emergency (overflow) occurred and after that time it was irrelevant.

Another example was the statement:

> "Interviews with operations personnel did not produce a clear reason why the response to the SO_2 alarm took 31 minutes. The only explanation was that there was not a sense of urgency since, in their experience, previous SO_2 alarms were attributed to minor releases that did not require a unit evacuation."

We are confused by the conclusion that there was no clear reason for the delayed response as the explanation provided in the second sentence seems completely clear and reasonable to us, particularly as nuisance alarms were relatively common (they occurred about once a month).[5] The delay in responding to the alarms seems understandable given the fact that past alarms had not been important, other (unrelated) alarms were sounding at the same time that required the operators attention, and all the information the operator had about lack of flow into the tank supported the conclusion that the tank level alarm was spurious and (after the analyzer alarm sounded) the release was minor and did not require a unit evacuation. The analyzer alarm maxed out at 25 ppm so the Board Operator had no indication of the real level (which was much higher).

To understand completely why the Board Operator's process model was so clearly incorrect and his behavior not what was, in hindsight, desirable, a full human-factors analysis of the operator interface to the plant is required. Such a human factors investigation should be part of all accident investigations where, as in this case, the operator was confused or unaware of the actual state of the plant or made assumptions that were incorrect. There is probably a reason for each of these mistakes and they need to be carefully investigated so that appropriate changes can be made for the future. There are very likely a lot of explanations for the operator's behavior related to system and plant interface design that a full human factors analysis would detect. We could not perform such an analysis due to lack of information about the design of the control room and the Board Operator's control interface.

There is one other clue provided in the accident report to explain some of the Board Operator's behavior. The report notes that the Board operator "did not demonstrate an awareness of risks associated with overflowing the tank and potential to generate high concentrations of SO_2 if the surfuric acid was spilled." There is no explanation of why this might have been true. Is there a deficiency in the training procedures about the hazards associated with his job responsibilities? If the explanation is that this particular operator is simply incompetent and although exposed to effective training did not profit from it, then the question becomes why such an operator was allowed to continue in that job and why the evaluation of his training outcomes did not detect this deficiency.

One of the common problems in systems with multiple controllers is confusion and coordination problems about who should be doing what and what the other people are doing. It is very likely that coordination problems occurred here, but we cannot tell from the information in the accident report alone. As will be noted later, the safety policy, as currently written, seems to ensure that such coordination problems will occur.

One other inadequate control action related to the loss was the lack of a work order reporting the erratic behavior of the level transmitter for a year and a half after it started acting erratically (January 2006). We have too little information to understand why this time lag occurred, but it is an important question that should have been answered in the accident report.

[5] This information was not in the accident report but was provided in a response to one of our questions Apparently such nuisance alarms are caused by sampling errors or other routine activities.

Outside (Field) Operator:

Safety-Related Responsibilities:

- Provide control actions that will avoid chemical release
- Sound evacuation alarm when an exposure occurs

Context:

Related to Risk and Procedures:

- No written unit procedure for responding to an SO_2 alarm. The "standard response" to an SO_2 alarm on the SBS unit is to have operator conduct a field assessment of the situation.
- No written procedure for ordering evacuation when an SO_2 alarm sounds or criteria established for the level of SO_2 that should trigger an evacuation alarm.
- Unit training material does contain information on the hazards of SO_2, including IDLH information, but this information has not been instituted in standard operating/emergency procedures.
- Time required to conduct a field assessment varies based on specific circumstances.
- Some said it was standard operating procedures to leave the manual block valves open while others said it was SOP to leave them closed. No written procedures or protocols provided

Related to Alarms:

- Distracted by other concerns related to transferring the Pit Sweep from the SBS to the Beavon-Stretford, which demanded his attention during the time of this incident. An alarm indicated that the RGG fire eye went out.
- No way to determine actual level of SO_2. Alarm maxes out at 25 ppm.
- Previous SO_2 alarms were attributed to minor releases that did not require a unit evacuation. Such alarms occur approximately once a month.
- In past, units were not evacuated by blowing the horn, but rather by operations personnel walking through the unit and stopping work so a "reluctance to hit the evaluation horn is apparent among some operations personnel."
- None of alarms designated as critical alarms. No established criteria in a procedure for what SO_2 levels constitute an emergency condition that should trigger sounding the evacuation alarm.

Control actions related to the loss

Control Valve:

- Leaves manual block valves in open position.
- At 10:00 Field Operator checks on the Beavon-Stretford unit at behest of board operator

Emergency Response:

- 31 minutes transpired from time SO_2 alarm sounded until Outside Operator reached unit, conducted an assessment, and implemented measures to stop the release.
- Put himself at risk by entering the immediate area to close the manual block valves instead of allowing properly equipped emergency personnel to handle the release.

Process Model Flaws:

Tank:

- Believed there was no flow into tank.
- Did not know flow meter was not working so thought tank was not filling when it was.
- Did not know that board operator had erratically functioning level indicators and a non-functioning alarm.

SO_2 Risk

- Knew a significant release was occurring but did not know analyzer had maxed out at 25 ppm or that IDLH concentration for SO_2 was 100 ppm.
- Was not aware of risk of SO_2 release via tank overflow.
- Did not know SO_2 release warranted evacuation order.

Emergency Response:

- Felt he had the authority to call for a unit evacuation but was not sure that conditions were bad enough to make that call during this incident.

Plant:

- Did not have "full awareness" of all work being performed on or near the unit

The process model for the Outside Operator could be used to understand his behavior as was done for the Board Operator. But his model is pretty simple: he believed there was no flow because the (non-functioning) flow meter said that was so. The Board Operator and Outside Operator both see the output from the same sensor indicating flow. In this case, that sensor was faulty. The only other information he had came from the Board Operator, whose process model and process feedback information was also incorrect. The Outside Operator has no other direct indicator of the level in the tank. He did not make a second effort to visually verify the flow at 9:25 because he was in a hurry to get to the simultaneous but unrelated trip of equipment in another part of the unit.

The primary mistake (in hindsight) made by the Outside Operator concerned the delay in the evacuation alarm and his attempt to clean up the spill instead of immediately seeking help. The official accident report says that the evacuation signal was delayed because the field operator was not sure the conditions were bad enough to make that call. The Outside Operator also seemed to have a poor understanding of the risks of an SO_2 release. Once again, several reasons could exist for this, but no matter what the reason, it is clearly something that needs to be remedied—not just for SO_2 but also for the risks of the plant in general. An audit should be performed to determine if SO_2 release is the only hazard that is not understood and if these two operators are the only ones who are confused.

In addition to examining the training and understanding of the hazards by the operations personnel, the procedures for sounding evacuation alarms need to be examined. If there is a potential for operators to make poor decisions in safety-critical situations, then they need to be provided with the criteria to make such a decision. Expecting operators under stress and perhaps with limited information and training to make such critical decisions based on their own judgment is unrealistic: It simply ensures that operators will be blamed when their decisions turn out, in hindsight, to be wrong.

The other training or procedural deficiency here beyond not sounding the alarm, i.e., the operator trying to fix the problem rather than calling in emergency personnel immediately, also

needs to be more carefully examined. In fact, this response is the *normal* one for humans; if it is not the desirable response then procedures and training must be used to ensure that a different response is elicited. The role of the current Company X safety policy here is discussed later.

If Company X wants the flexibility inherent in real-time decision making rather than strict procedures, then they will need to provide much more extensive training and better real-time information to the operators.

HSE (Process Safety):

Safety-Related Responsibilities:

- Perform hazard analyses and identify safety-critical procedures and equipment
- Conduct risk assessments and develop operating procedures that adequately control the risk associated with safety-critical operations.
- Inform operators about critical processes and alarms.
- Audit safety-related equipment
- Provide MoC and CoW procedures for safety-critical operations.
- Ensure that all plant personnel have adequate training about:
  - The risks of chemicals and materials at their place of work
  - The risks associated with their job
  - How to perform their job safely
  - Proper emergency response

Context: Unknown

Control actions related to the loss

- Unit training information does contain information on the hazards of SO_2 but this information has not been instituted in standard operating/emergency procedures.
- None of the high level alarms were designated as critical alarms
- Did not provide any risk-assessed operating procedure for drawing down acid and filling Tank 731.
- No specific unit procedures or other protocols that define critical operational parameters were provided
  - Sequence to initiate the drawdown process (e.g., notification of Outside Operator, manual block valve positioning, etc.)
  - Process control parameters (e.g. drawdown initiation and endpoint, specified flow rate into the tank, safe level at which Tank 731 is considered full)
  - Sequence of steps necessary to conclude and secure the tank filling process (e.g., closing block valves)
  - Appropriate response to alarms (e.g. high level alarm, SO_2 alarm)
  - Clear conditions for evacuation, including SO_2 ppm level "There is not established criteria in a procedure for what SO_2 levels and/or alarms constitutes an emergency condition which should trigger sounding the evacuation alarm"

Process Model Flaws: Unknown

Because the report focused on the operators' role in the incident and not on the role of other groups, we cannot provide information about the context or the process model flaws that might have existed for components of the safety control structure that were involved in the accident

process but were not examined in the accident report. The inadequate control actions included on the previous page did occur, but there is no way to determine the reason for them. This would be important information to obtain in determining how to fix any systemic problems. For example, how are critical alarms identified and why were these alarms not identified as such when they clearly were involved in a serious incident? Is such a determination done using HAZOP? If so, are there problems in performing HAZOP effectively in these circumstances? Alternatively, are the problems perhaps in the criteria used to label alarms as critical?

We asked whether the HAZOP had identified the alarms as critical and were told that it had not identified a critical alarm for this application. A hypothesis for why not is that it was "probably because it was an infrequent batch operation." If HAZOP is not useful for infrequent batch operations, then a different hazard analysis method should be used or the HAZOP procedures changed or augmented.

We were told that the omission of criteria for which SO_2 levels and/or alarms constitutes an emergency that should trigger sounding the evacuation alarm was simply an oversight. What procedures are in place to detect or prevent such oversights? It seems that any identified hazards should have procedures in place for responding if the hazard occurs so either the hazard was not identified or there was no procedure for checking that there is a response for all identified hazards.

What criteria have been established at Company X for determining what emergency procedures are provided to the operators? The report says that there is a safety policy (C-13) that states: *"At units, any employee shall assess the situation and determine what level of evacuation and what equipment shutdown is necessary to ensure the safety of all personnel, mitigate the environmental impact and potential for equipment/property damage. When in doubt, evacuate."* There are two problems with such a policy.

The first problem is that evacuation responsibility does not seem to be assigned to anyone but can be initiated by all employees. The consequence of such lack of assigned control responsibility is usually that everyone thinks someone else will take the initiative. Responsibility for sounding alarms must be assigned to someone. Others may report problems and even sound an evacuation alert when necessary, but there must be someone who has the actual responsibility, accountability, and authority, and there should be backup procedures for others to step in when that person does not execute their responsibility acceptably.

The second problem is that unless the procedures clearly say to sound an alarm, humans are very likely to try to diagnose the situation first. The same problem pops up in many accident reports—the operator is overwhelmed with information that they cannot digest quickly or do not understand and will first try to understand what is going on before they sound an alarm. If Company X wants employees to sound alarms expeditiously and consistently, then the safety policy should specify exactly when alarms are required, not leave it up to personnel to "evaluate a situation" in which they probably are confused and unsure about what is going on (as in this case) and under pressure to make quick decisions under stressful situations. How many people, instead of dialing 911 immediately, try to put out a small kitchen fire themselves? That it often works simply reinforces the tendency to act in the same way during the next emergency. And it avoids the embarrassment of the firemen arriving for a non-emergency.

Maintenance:

Safety-Related Responsibilities:

- Ensure that all safety-related equipment is in proper working order.

Context:

- Large backlog of maintenance work orders associated with unit equipment.

Control actions related to the loss

- Did not meet target of 4 weeks (28 days) average age of a "safety" work order. Average age was 86 days. In particular, did not fix flow meter (FT47706) and level transmitter (L47731A) involved in the incident for an extended time.
- The level transmitter L47731A had supposedly been fixed on July 25, but it was not functioning properly at the time of the incident 2 weeks later. Was this a result of inadequate repair or perhaps inadequate testing after the repair?

Process Model Flaws:

Unknown

Questions that need to be answered here are why the level transmitter was not working so soon after it was supposedly fixed, why safety work orders were so delayed, whether plant and corporate management knew this and what they were doing about it (if anything), etc. If the targets are impractical and unobtainable, then the targets need to be re-evaluated or more resources applied. There was a seemingly large number of non-functioning or erratically functioning equipment in this unit. Was this unusual or is it a problem throughout the refinery (and, perhaps, throughout the company)? Is the problem actually in HSSE because critical alarms are not being properly identified?

In response to our questions, we were told that the Maintenance Manager was aware of the backlogs and was working on reducing them. Some of the problem was bad data—work orders that had been completed but not closed, work orders marked as safety related but not really a safety issue, etc. Because of the difficulty in managing something with bad data, ways to improve the data should be developed.

Operations Management:

Safety-Related Responsibilities:

- Ensure operators have the information they need to carry out their tasks, including information about the current state of the equipment they are operating
- Ensure operators have been given adequate operational procedures and training on how to implement them and audit to ensure that the training is effective.
- Ensure operators understand the hazards associated with the equipment they are operating and audit to ensure that training is effective.
- Provide reporting channels for malfunctions and assign responsibility. Ensure that communication channels are effective and operating correctly.
- Ensure emergency procedures are being executed properly.

Context: Unknown

Control actions related to the loss:

- Allowed unreported malfunctioning safety-related equipment for 2.5 years
- Did not respond properly when past evacuation alarms were provided inadequately (no horn and use of verbal commands only)
- Was appropriate training provided? Were there audits to check the effectiveness of the training? Gun drills were apparently not run on this process.

Process Model Flaws: Unknown

Again, it is difficult to analyze the role of operations management because so much was left out of the accident report.

Emergency procedures are the basis for gun drills at the refinery. But there were no emergency procedures provided for this procedure. We learned that there is a refinery "gun drill" policy and matrix that list requirements (which procedures, how often, etc.). This policy needs to be reviewed as adequate gun drills were not done for this unit. The Company X accident report recommends running a gun drill on this unit but not evaluating the policy for determining when gun drills are done.

The training coordinator is responsible for training operators on hazards, and we have been told that unit training includes a review of the hazards of the acid stream. The effectiveness of this training should be evaluated given the incident and the operators' responses.

The Control of Work status board was not effective in notifying the operators about where work was occurring around the unit. Why not?

Why the evacuation horn had been delayed or improperly used in the past has not been thoroughly investigated. Two possible reasons provided to us in response to our questions is that (1) the hazard or extent of the situation was not understood and/or (2) the operators tried to handle the situation internal to the unit instead of sounding the evacuation alarm (which calls attention to the incident). This should be investigated further beyond the reasons hypothesized in the report that a precedent had been set. While that may be true, there are many other even more plausible reasons that could have caused it. Solutions can then be proposed. It is also important

to understand why management did not know about this behavior or, if they did, why they did not take measures to change it. It seems extremely important to answer such questions.

Plant Management:

Safety-Related Responsibilities:

- Ensure that plant is operated in a safe manner, that appropriate safety responsibilities (and authority and accountability) have been assigned, and that everyone is carrying out their safety-related responsibilities.

Context:

Unknown, probably lots of performance pressures together with less pressure for his or her safety-related responsibilities.

Control actions related to the loss

- Allowed a large maintenance backup of safety-related equipment over the target dates
- Allowed lots of safety lapses in plant operations.

Process Model Flaws:

- Probably lots of wrong information, e.g., did they know about maintenance backlogs? Did they know about lack of risk-assessed procedures? Did they understand the risk in current plant operations?

Plant management has responsibility for ensuring that the plant is run safely. Either the plant management did not know about the inadequacies of operations at the plant (and that is the problem that must be fixed) or they did not take adequate steps to fix the problems. The latter could have resulted from many factors, none of which are included in the accident report so we cannot determine why the necessary steps were not taken. Some common reasons are that the feedback about safety at the plant is inadequate, safety risk is judged to be low and other risks (financial, etc.) are judged to be greater, reward structures for management are not appropriate, etc. In a complete accident/incident report, the role of plant management must be investigated and analyzed as carefully as that of the operators and lower-level managers.

Company X Corporate Management:

Safety-Related Responsibilities:

- Ensure that company facilities are operated in a safe manner, that appropriate safety responsibilities (and authority and accountability) have been assigned, and that plant managers are carrying out their safety-related responsibilities.
- Establish an appropriate reward structure for safe behavior.

Context:

Unknown, probably lots of performance pressures and less pressure for safety-related responsibilities and oversight

Control actions related to the loss

- Allowed at least one refinery to be operated with many safety deficiencies.

Process Model Flaws:

- Probably lots of wrong information, e.g., did they know about maintenance backlogs? Did they know about lack of risk-assessed procedures?

Again, little information is available. And, again, the reasons for the inadequate control over safety at this level needs to be identified and fixed.

Corporate HSE:

Safety-Related Responsibilities:

- Ensure that plant is operated in a safe manner, that appropriate safety responsibilities have been assigned, and that everyone is carrying out their safety-related responsibilities.

Context:

Unknown, probably lots of performance pressures and less pressure and for safety-related responsibilities and oversight

Control actions related to the loss

- Did not provide high-level policies for performing effective hazard analysis, etc. or they were provided but did not know about the deficiencies in their execution.

Process Model Flaws:

- Probably lots of wrong information, e.g., did they know about maintenance backlogs? Did they know about lack of risk-assessed procedures?

Once again, there is no information about the role of this group in the incident, but from a systems thinking perspective, they clearly played a role. We will not speculate about the reasons for the inadequate control at this level.

Communication and Coordination

There was inadequate communication from the operations to contractors working in the area that an incident involving release of sulfuric acid had occurred and that the SO_2 alarms had sounded. The communications between the Board and Outside operators seem good. Not enough information was provided to evaluate other types of critical communications within the plant and within Company X.

Dynamics and Changes over Time

We have little information about the dynamic aspects of safety that might have contributed to the T-731 overflow incident. But the answers to some questions we asked gave hints that changes over time may have been involved in the problems. For example, to our question about why there was no established criteria in a procedure for what SO_2 levels and/or alarms constitutes an emergency that should trigger sounding the evacuation alarm, the answer provided was "Just not thought of before—perhaps it was 'not required' before when there were many experienced personnel in the units." Has the level of experience of personnel changed significantly over time? In general, has the behavior of the components in the safety control structure degraded and why?

Recommendations

Equipment:

Not having access to engineering design information for the plant, it is difficult for us to recommend physical design changes not recommended in the original report. For example, we do not know why the sewer was blocked or whether a redesign might eliminate that cause. Nevertheless, three equipment recommendations arise from the STAMP analysis that were absent from the Company X accident report:

1. Change the SO_2 alarm to be audible and can be heard by people in the affected areas who are not in the line of sight of the strobe light alarm.
2. Determine why the sewer was blocked and improve the design or procedures to avoid this in the future.
3. Safety-critical indicators should show actual state and not just commanded position. A hazard analysis could be used to determine where additions or changes are needed.

Management and Engineering:

1. Review procedures for deciding when to do risk assessments of operating procedures.
2. Establish and enforce procedures for writing work orders for non-routine operations.
3. Evaluate normal operating levels of tanks and alarm settings (is it appropriate to routinely operate a tank at 0.3 ft below the alarm threshold?).
4. Investigate frequent spurious alarms. Determine the root causes and establish ways to reduce their frequency.
5. Consider adding an alarm to indicate rising tank level when the control valve has been commanded closed.
6. Consider a way to easily and clearly mark malfunctioning gauges/meters for the operators.
7. Consider adding a policy where manual block valves are left closed when it appears that flow is not occurring and should not occur instead of depending on the Board Operator's control valve only (creates a single point of failure).
8. If the standard response to an alarm is to have a field operator check it, then need some instructions on what to do if the field operator is not immediately available. Alternatively or in addition, use dedicated operators for non-routine batch operations (as this was) when practical.
9. Define safe operating limits for all safety-related equipment. Establish corporate policy for doing this and ensure it is being followed.
10. Improve malfunction reporting and communication, e.g., establish a way to communicate information about all current malfunctioning equipment to operators. Determine if there are any barriers to reporting malfunctioning equipment and, if so, eliminate them.
11. Clarify the evacuation alarm procedures. Company X should not just leave it up to the operators to assess the situation, particularly when the operator does not have all the necessary information to make a proper decision. Rewrite the general safety policy (C-13) on evacuation procedures. [This recommendation conflicts directly with the related recommendation in the Company X accident report, which suggests reinforcing the current policy with the operators.]
12. Investigate why in the past the evacuation alarm was not handled appropriately.
13. Perform a full human-factors analysis of the operator's interfaces to the plant to determine what changes are necessary to ensure improved control by the operators in the

future. This will require a major effort but will be much less effort if such analyses are done (as they should be) at the time the interfaces are originally designed.

14. Redesign maintenance activities or provide additional resources to ensure critical work orders are completed on time. If the targets are impractical and unobtainable, then evaluate the targets and the risk involved in changing them to more realistic values. Improve the procedures for collecting data on target achievement if that is the problem.
15. Create a system to ensure that operators understand the hazards associated with the processes they are controlling. Evaluate their knowledge and revise training procedures if necessary.
16. Evaluate and perhaps improve the CoW procedures with respect to knowing what personnel (contractors) are in the area. Determine how to provide more effective "situation awareness" of all work being performed on or near the units.
17. Evaluate the information operators have about the plant hazards and ensure they have what is required to operate hazardous processes safely both in terms of the design of feedback and in terms of ensuring that feedback channels are operational and working properly.
18. Evaluate current policy with respect to operating processes with safety-critical equipment that is known to be non-operational and establish policy if it does not already exist. While it may at times be necessary, there should be an approval process by someone who has the proper authority to evaluate and accept the risks and information should be provided to corporate management about when and how often this occurs. Before safety-critical, non-routine, potentially hazardous batch operations are conducted, safety policies should require inspection of all safety equipment involved to ensure it is operational, including testing of alarms.
19. Evaluate all operating procedures to determine whether they need to be risk-assessed and, even more important, determine how this will be done for new or changed procedures in the future. This is clearly a very large task and will take time to finish.
20. Evaluate the refinery gun-drill policy.
21. Improve the methods used to identify hazards, critical alarms and parameter values for all hazards (including non-routine batch operations), design feedback to provide operators with the real-time information they need to deal with hazards, and create procedures for operators to deal with all hazards.
22. Evaluate the safety control structure at the refinery and at the Company X Corporate level to determine whether it is adequate to ensure safety and make improvements if it is not. Ensure explicit specification of each component's responsibilities (and authority and accountability) and make sure each person assigned the responsibilities knows what they are and how to execute them. Determine whether each controller has the information required to effectively carry out their safety-related responsibilities. Establish procedures to periodically audit feedback channels to ensure they are operating correctly.

3.4 Comparison with Comprehensive List of Causes

Company X trains accident investigators to use an approach based on the Swiss Cheese Model and causal factor checklists to identify and analyze causes. The Baker Panel report was very critical of this type of checklist approach to accident analysis. A major concern stated in the panel report is that using a checklist will lead to not identifying any factors other than those on the list. In the panel's experience, investigators typically use a checklist as a complete list of potential causes instead of as a starting point for a discussion of the deeper systemic causes and usually will not identify factors that are not on the list. We found that to be true in the Company X refinery report. Even though causal factors were identified in the narrative of the report, they did not lead to recommendations when they did not correspond to items in the Company X checklist.

The official accident report identified eight recommendations. These eight recommendations, all of which were linked to parts of the checklist (which we assume is established practice for the company and thus belies the belief that investigators will go beyond the checklist) are listed below. We compare them with the recommendations we generated using STAMP-based analysis of the accident.

1. Operator duty to respond to alarms needs to be reinforced with the work force.
   [This recommendation ignores the reasons for why the operators did not respond to the alarms and changing the contextual and behavior-shaping factor in the system design.]
2. Two alarm points (high and high/high) should be established for each of the redundant level sensors on Tank 731. The high alarm should be designated as a critical alarm given the risks associated with overflowing this tank.
   [Reasonable, but it ignores the other alarms at the plant. Could there be other incorrectly identified alarms and is there a flaw in the procedures for identifying critical alarms? While necessary, this action will not solve the whole problem. The alarm was ignored because it was believed to be in error, not because it was believed to be non-critical and therefore acceptable.]
3. Consideration should be given to establishing two alarm set-points for the SO_2 analyzers. The current alarm setpoint at 4 ppm provides an important warning of the release of SO_2 gas that should be investigated. Consideration should be given to establishing a second alarm setpoint at a level that triggers an immediate evacuation alarm (e.g., when the instrument maximum reading is reached).
   [Good.]
4. Implement new control of work procedures to achieve better operational awareness and control of work being performed on the unit. Process hazards and potential risks associated with operational activities should be identified by the AA and communicated to the PA for inclusion in risk assessments and crew reviews to improve the awareness of all work groups about the area hazards. Utilize process risk assessment to make more informed decisions about when to best schedule potentially hazardous operations to minimize potential risk to work crews.
   [Good]
5. Develop a risk assessed procedure for the acid drawdown process in accordance with Policy D-22 which defines critical operational parameters such as the sequence of steps required to initiate the drawdown process (e.g., notification of Outside Operator, manual block valve positioning, etc), process control parameters (e.g., drawdown initiation and endpoint,

specified flow rate into the tank, etc.), the safe level at which Tank 731 is considered full, the sequence of steps necessary to conclude and secure the tank filling process (e.g., closing block valves), and appropriate response to alarms.

[Why just this process?]

6. Consideration should be given to conducting a gun drill on the unit with a focus on how to recognize and report emergency response conditions, proper communications, and the circumstances under which a unit evacuation should be conducted.

   [Again, we believe this recommendation should be generalized to evaluating the overall gun drill policy.]

7. Unit evacuation procedure should be revised and emphasize that employees shall assess the situation and determine what level of evacuation is necessary to ensure the safety of all personnel as stated in Safety Policy C-13. Consideration must include ongoing SIMOPs (as posted on CoW boards) in determining evacuation requirements.

   [We disagree with this recommendation as stated earlier.]

8. Convene a cross function team of refinery and OCC personnel to examine "protection of workforce" issues. Areas of focus should include:
   - Evacuation – understanding and awareness of workforce
   - Evacuation – when to activate

   Specific issues to consider include:
   - Identification of potential conditions that could lead to loss of containment by the process unit.
   - Identification and planning to evacuate during unit upsets and transients, loss of instruments/view, exceeding safe operating limits / critical alarms, loss of level or filling a drum / tower, etc.

   Team recommendations should apply site-wide.

   [A good recommendation, but it could be more general. For example, what about the effectiveness of the alarms in evacuation (visual vs. aural) and the reasons why evacuation alarm procedures have not been executed correctly in the past.]

Our analysis of the T-731 accident and recommendations were more general and less focused on this particular incident and fixing the specific problems of this particular refinery in the company. While several of the Company X report recommendations do tackle systemic factors, most are limited strictly to the details of this particular incident. The first one, i.e., to reinforce operator duty to respond to alarms, results from the use of hindsight bias in blaming the operators for not responding without considering the systemic factors that led them to respond inappropriately. Simply telling the operators to respond to alarms without changing those factors will not be very effective.

We were limited to the information contained in the official accident report and could not evaluate why the flawed management control actions and decision making related to the incident occurred at the higher levels of the safety control structure. To do this, we would need more information about the context in which these actions took place and about the information provided by the feedback channels to the controllers.

In summary, while the basic analysis in the Company X accident report was fine, as far as it went, we were surprised by the limited nature of the recommendations and by the examples of hindsight bias we found. In general, the report focused on operator error, as does the Company X causal factors checklist, even in some of the descriptions of the identified "system causes." Such a limited focus will limit the usefulness of accident reports in providing fixes to prevent future accidents.

4. Conclusions

In this research project, we demonstrated the application of systems thinking to the analysis of an accident at a Company X refinery and compared the results with the official report.

Our experience is that the official refinery accident report is representative of most accident reports at Company X. The latest version of their checklist for accident analysis was used in their analysis of the spill and the report was written by a trained and competent group of accident investigators and experts. We believe the problem is the approach to accident analysis required in Company X, not the investigators.

There is good news and bad news here. The good news is that we used only the information we found in the official accident report so Company X could learn much more from adverse events without much extra effort in terms of collecting information. The bad news is that continuing to emphasize old approaches to thinking about accidents can serious hamper the company's efforts to improve safety. Rather than expanding the investigator's analysis, the checklist seemed to limit their conclusions and forced their thinking down narrow paths. The current training materials on accident analysis at Company X also emphasize blaming operator error and not on understanding why the particular operator errors occurred. Together these two limitations can lead to missing most of the systemic causal factors and recommendations that could provide the most effective ways to prevent future accidents.

The next step would be to produce a tool to assist in an accident analysis using STAMP and in visualizing the causal analysis of the accident. We are hoping to get funding to be able to pursue this line of research.

5. References

1. Ayres, Robert U. and Rohatgi, Pradeep (1987), Bhopal: Lessons for technological decision-makers, *Technology in Society*, 9:19-45.
2. Bogard, William (1989), *The Bhopal Tragedy*, Westview Press, Boulder, Colorado.
3. Company X (2002). "Getting HSE right: A guide for Company X managers," December.
4. Checkland, Peter (1981), *Systems Thinking, Systems Practice*, John Wiley and Sons, New York.
5. Chisti, Agnees (1986), *Dateline Bhopal*, Concept Publishing Company, New Delhi.
6. Dekker, Sidney (2007), *Just Culture: Balancing Safety and Accountability*, Ashgate.
7. Dekker, Sidney (2009), Report of the Flight Crew Human Factors Investigation of the TK1951 Boeing 737-800 Accident near Amsterdam Schipol Airport Feb. 25 2009, Lund University School of Aviation, Sweden.
8. Ladd, John (1987), Bhopal: An essay on moral responsibility and civic virtue, Dept. of Philosophy, Brown University, January.
9. Leveson, Nancy (2003), A new accident model for engineering safer systems, *Safety Science*, 42(4), April, pp. 237-270.
10. Leveson, Nancy (2007), Technical and managerial factors in the NASA Challenger and Columbia losses: Looking forward to the future, in Handelsman and Fleishman (editors), *Controversies in Science and Technology: From Chromosomes to the Cosmos*, Vol. 2, Mary Ann Liebert Inc.
11. Leveson, Nancy (2012), *Engineering a Safer World*, MIT Press (electronic version available for download at http://sunnyday.mit.edu/safer-world)

12. Perrow, Charles (1986), The habit of courting disaster, *The Nation*, pp. 346-356, October.
13. Rasmussen, Jens (1997), Risk Management in a Dynamic Society: A Modelling Problem, Safety Science, Vol. 27, No. 2/3, Elsevier Science Ltd., pp. 183-213.
14. Reason, James (1990), *Human Error*, New York: Cambridge University Press.