CAST HANDBOOK:

How to Learn More from Incidents and Accidents

Nancy G. Leveson

An accident where innocent people are killed is tragic,
but not nearly as tragic as not learning from it.

Preface

About 15 years ago, I was visiting a large oil refinery while investigating a major accident in another refinery owned by the same company. The head of the safety engineering group asked me how they could decide which incidents and accidents to investigate when they had hundreds of them every year. I replied that I thought he was asking the wrong question: If they investigated a few of them in greater depth, they would not have hundreds. I don't think he understood, or at least did not accept, my suggestion. The goal of this handbook is to explain that answer—we are not learning enough from the incidents and accidents we are having. We need to figure out how to learn more if we truly want to significantly reduce losses.

After working in the field of system safety and helping to write the accident reports of several major accidents (such as the Space Shuttle Columbia, Deepwater Horizon, and Texas City) and other smaller ones, I have found many factors common to all accidents. Surprisingly, these are often not included as a cause in the official accident reports. CAST (Causal Analysis based on System Theory) and this handbook are my attempt to use my experience to help others learn more from accidents in order to do a better job in preventing losses in the future.

The handbook describes a structured approach, called CAST (Causal Analysis based on System Theory), to identify the questions that need to be asked during an accident investigation and determine why the accident occurred. CAST is very different than most current approaches to accident analysis in that it does not attempt to assign blame. The analysis goal changes from the typical search for failures to instead look for why the systems and structures in place to prevent the events were not successful. Recommendations focus on strengthening these prevention (control) structures, based on what was learned in the investigation.

How best to perform CAST has evolved with my experience in doing these analyses on real accidents. Updates to this handbook will provide more techniques as all of us learn more about this systems approach to accident analysis.

Acknowledgements:
I would like to thank several people who helped to edit this handbook: Dr. John Thomas, Andrew McGregor, Shem Malmquist, Diogo Castilho, and Darren Straker.

TABLE OF CONTENTS

TABLE OF FIGURES

Chapter 1: Introduction

My goal for this handbook is not to provide a cookbook step-by-step process that you can follow like a recipe. While that is often what people want, the truth is that the best results are not obtained this way. Instead, they are generated by providing ways for experts to think carefully and in depth about the cause of an accident. We need tools that are able to encourage broader and deeper thinking about causes than is usually done. In this way, it is my hope that we are able to learn more from events.

It is always possible to superficially investigate an accident and not learn much of anything from the effort. The same accidents then occur over and over and are followed each time by the same superficial analyses. The goal instead should be to invest the time and effort needed to learn enough from each accident so that losses are dramatically reduced and fewer investigations are required in the future.

Why do we need a new accident analysis tool?

The bottom line is that we are learning less from losses and near misses than we could. There are many accident analysis tools that have been created, particularly by academics, but few have significantly reduced accidents in real systems or even been used widely. Most focus on new notations for documenting the same old things.

If you want an academic and philosophical treatment of the subject, I recommend my "learning from events" paper[1] and my book *Engineering a Safer World* [Leveson 2012]. Reading *Engineering a Safer World* will help you to more deeply understand the limitations of current accident analysis approaches and assumptions and the technical and philosophical underpinnings of CAST. But that is not the goal of this handbook.

Instead, the goal here is to provide a practical set of steps to help investigators and analysts improve accident reports. Accident investigations too often miss the most important causes of an accident, instead choosing to focus on only one or two factors, usually operator error. This oversimplification of causality results in repetitions of the same accident but with different people involved. Because the *symptoms* of each loss seem to differ, we fix those symptoms but not the common underlying causes. As a result, we get stuck in continual fire-fighting mode.

What you will learn

This handbook will teach you how to get more useful results from accident investigation and analysis. While it may be necessary to spend more time on the first few accident analyses using this approach, most of the effort spent in modeling and analysis in your first use of CAST will be reused in subsequent investigations. Over a short time, the amount of effort should be significantly reduced with a net long term gain not only in a reduction in time spent investigating future accidents but also in a reduction of accidents and thus investigations. Experienced accident investigators have found that CAST allows them to work faster on the analysis as it creates the questions to ask early, preventing have to go back later.

Your long-term goal should be to increase the overall effectiveness of the controls used to prevent accidents. These controls are often embedded in a Safety Management System (SMS). Investigating accidents and applying the lessons learned is a critical part of any effective SMS. In turn, the current weaknesses in your SMS itself will be identified through a thorough accident/incident analysis process. Investing in this process provides an enormous return on investment. In contrast, superficial analysis of

[1] Nancy Leveson, Applying Systems Thinking to Analyze and Learn from Events, Safety Science, Vol. 49, Issue 1, Januagey 2010, pp. 55-64.

why accidents are occurring in your organization or industry will primarily be a waste of resources and have little impact on future events.

In fact, the systemic causes of accidents even in diverse industries tend to be remarkably similar. In my career, I have been involved in the investigation and causal analysis of accidents in aviation, oil and gas production, space, and other fields as well as studying hundreds of accident reports in these and in most every other industry. The basic causal factors are remarkably similar across accidents and even industries although the symptoms may be very different. The types of omissions and oversimplifications in accident reports also tend to be very similar. That's actually good news because it means that there are lots of opportunities to improve learning from the past if we have the desire and the tools to do so. Sharing the results from CAST analyses that identify common systemic causes of losses will allow us to learn from others without having to suffer losses ourselves.

The STPA Handbook [Leveson and Thomas, 2018] teaches how to prevent accidents before they occur, including how to create an effective safety management system. But there are still likely to be accidents or at least near misses that occur, and sophisticated and comprehensive accident/incident analysis is an important component of any loss prevention program. With the exception of the U.S. Nuclear Navy program called SUBSAFE (described in Chapter 14 of *Engineering a Safer World*), no safety programs have eliminated all accidents for a significant amount of time. SUBSAFE has some unique features in that it severely limits the types of hazards considered (i.e., submarine hull damage leading to inability to surface and return to port), operates in a restricted and tightly controlled domain, and spends significant amounts of resources and effort in preventing backsliding and other factors that increase risk over time.

But even if one creates a perfect loss prevention program, the world is continually changing. While some changes are intentional and can be controlled, others are unplanned—systems themselves change over time, people's behavior within a system changes, and the environment within which the system operates will also change. Detecting the unsafe changes, hopefully by examining leading indicators of increasing risk (see Chapter 6 of the STPA Handbook) and thoroughly investigating near-misses and incidents using CAST, will allow unplanned changes to be identified and addressed before losses result.

There is no set notation or format provided in this handbook that must be used, although some suggestions are provided. The causes of different accidents may be best explained and understood in different ways. The content of the results, however, should not differ. The goal of this handbook is to describe a process for thinking about causation that will lead to more comprehensive and useful results. Those applying these ideas can create formats to present the results that are most effective for their own goals and their industry.

What is CAST?

The causal analysis approach taught in this handbook is called CAST (Causal Analysis based on System Theory). Like STPA [Leveson 2012, Leveson and Thomas 2018], the loss involved need not be loss of life or a typical safety or security incident. In fact, it can (and has been) used to understand the cause of any adverse or undesired event that leads to a loss that stakeholders wish to avoid in the future. Examples are financial loss, environmental pollution, mission loss, damage to company reputation, and basically any consequence that can justify the investment of resources to avoid. The lessons learned can be used to make changes that can prevent future losses from the same or similar causes.

Because the ultimate goal is to learn how to avoid losses in the future, the causes identified should not be reduced to an arbitrary "root cause." Instead, the goal is to learn as much from every accident as possible. This goal is what CAST is designed to achieve. Some accident investigators have actually complained that CAST creates too much information about the causes of a loss. But, is a simple explanation your ultimate goal? Or should we instead be attempting to learn as much as possible from

every causal analysis? Learning one lesson at a time and continuing to suffer losses each time is not a reasonable course of action. Systemic factors are often omitted from accident reports, with the result that some of the most important and far reaching causes are ignored and never fixed. Saving time and money in investigating accidents by limiting or oversimplifying the causes identified is false economy. The concept of "root cause" and "probable cause" are common ways to find someone or something to blame and then get on with life and business—until the next accident. The overriding question is whether to pay now or pay later.

Relationship Between CAST and STPA

Hazard analysis has been described as "investigating an accident before it occurs." STPA (System Theoretic Process Analysis) is a hazard analysis tool based on the same powerful model of causality as CAST. In contrast to CAST, its proactive analysis can identify all potential scenarios that may lead to losses, not just the scenario that occurred. These potential scenarios produced by STPA can then be used to prevent accidents before they happen. CAST, in contrast, assists in identifying only the particular scenario that occurred. Although their purposes are different, they are obviously closely related. Because STPA can be used early in the concept development stage of an accident (before a design is created), it can be used to design safety and security into a system from the very beginning, greatly decreasing the cost of designing safe and secure systems: Finding potential safety and security flaws late in the design and implementation can significantly increase development costs. CAST analyses of past accidents can assist in the STPA process by identifying plausible scenarios that need to be eliminated or controlled to prevent further losses.

Format and Use of this Handbook

This handbook starts with a short explanation of why we are not learning as much from accidents as we could be. Then the goals and the process for performing a CAST analysis are described. A real example of a chemical plant explosion in the Netherlands is used throughout. The causal factors in this accident are similar to most accidents. Many other examples of CAST analyses can be found in *Engineering a Safer World* and on the PSAS website (http://psas.scripts.mit.edu). Appendix A provides links to CAST analyses in a wide variety of industries.

The worlds of engineering safety and workplace safety tend to be unnecessarily separated with respect to both the people involved and the approaches used to increase safety. In fact, this separation is unnecessary and is inhibiting improvement of workplace safety. A chapter is included in this handbook on how to apply CAST to workplace (personal) safety.

While CAST and structured accident analysis methods have been primarily proposed for and applied to accidents involving physical systems, CAST can very effectively be used on social system "accidents," which may entail major disruptions, loss of life, or financial system losses. Examples are shown in Chapter 5 for a pain management drug (Vioxx) that led to serious physical harm before being withdrawn from the market and for the Bears Stearns investment bank failure in the 2008 financial system meltdown.

In summary, while there are published examples of the use of CAST as well as philosophical treatises on the underlying foundation, there are presently no detailed explanations and hints about how to do a CAST analysis. The goal of this handbook is to fill that void.

CAST is based on fundamental engineering concepts. For readers who do not have an engineering background, Appendix E will provide the information necessary to understand this handbook and perform a CAST analysis.

Chapter 2: Starting with some Basic Terminology

"When I use a word," Humpty Dumpty said, in rather a scornful tone, "it means just what I choose it to mean—neither more nor less." "The question is," said Alice, "whether you can make ***words*** *mean so many different things." "The question is," said Humpty Dumpty, "which is to be master—that's all."*

Lewis Carroll (Charles L. Dodgson), *Through the Looking-Glass,* first published in 1872.

While starting from definitions is a rather dull way to start talking about an important and quite exciting topic, communication is often inhibited by the different definitions of common words that have developed in different industries and groups. Never fear, though, only a few common terms are needed, and this chapter is quite short.

As Humpty Dumpty (actually Charles Dodgson) aptly put it, the definitions established here apply to the use of this handbook, but are not an attempt to change the world. There is just no way to communicate without a common vocabulary.

*Accident (*sometimes called a *Mishap)*: An undesired, unacceptable, and unplanned event that results in a loss. For short, simply a loss.

Undesirability and unacceptability must be determined by the system stakeholders. Because there may be many stakeholders, a loss event will be labeled an accident or mishap if it is undesirable or unacceptable to any of the stakeholders. Those who find the loss desirable and acceptable will not be interested in preventing it anyway so to them this book will be irrelevant.

Note that the definition is extremely general. Some industries and organizations define an accident much more narrowly. For example, an accident may be defined as only related to death of or injury to a human. Others may include loss of equipment or property. Most stop there. The definition above, however, can include any events that the stakeholders agree to include. For example, the loss may involve mission loss, environmental pollution, negative business impact (such as damage to reputation), product launch delays, legal entanglements, etc. The benefit of a very broad definition is that larger classes of problems can be tackled. The approach to accident analysis described in this book can be applied to analyzing the cause of any type of loss.

It is also important to notice that there is nothing in the definition that limits the events to be inadvertent. They may be intentional so safety and security are both included in the definition. As an example, consider a nuclear power plant where the events include a human operator or automated controller opening a valve under conditions where opening it leads to a loss. The loss is the same whether the action was intentional or unintentional, and CAST can be used to determine why it occurred.

Universal applicability of the accident definition above is derived from the basic concepts of system goals and system constraints. The system goals stem from the basic reason the system was created: such as producing chemicals, transporting passengers or cargo, waging warfare, curing disease, etc. The system constraints are defined to be the acceptable ways those goals can be achieved. For example, it is usually not acceptable to injure the passengers in a transportation system while moving them from place to place. Events that damage the company's reputation while achieving short term profits may also not be acceptable to the stakeholders.

To summarize:

System Goals: the reason the system was created in the first place

System Constraints: the ways that the goals can acceptably be achieved

Notice here that the constraints may conflict with the goals. An important first step in system engineering is to identify the goals and constraints and the acceptable tradeoffs to be used in decision making about system design and operation. Using these definitions, system reliability is clearly not synonymous with system safety or security. A system may reliably achieve its goals while at the same time be unsafe or insecure or vice versa. For example, a chemical plant may produce chemicals while at the same time release toxins that pollute the area around it and harm humans. These definitions also explain why reliability may, in some cases, conflict with safety. A statement that a system "failed" does not provide enough information to understand what occurred or what goals or constraints were violated.

Two more definitions are needed. One is straightforward while the other is a little more complicated. The first is the definition of an incident or near-miss.

> *Incident* or *Near-Miss*: An undesired, unacceptable, and unplanned event that does not result in a loss, but could have under different conditions or in a different environment.

The final term that needs to be defined and used in CAST is *hazard* or *vulnerability*. The former is used in safety while the latter in security but they basically mean the same thing. A vulnerability is defined as a flaw in a system that can leave it open to attack while, informally, a hazard is a state of the system that can lead to an accident or loss. More formally and carefully defined:

> *Hazard or vulnerability*: A system state or set of conditions that, together with specific environmental conditions, can lead to an accident or loss.

As an example, a hazard might be an aircraft without sufficient propulsion to keep it airborne or a chemical plant that is releasing chemicals into the environment. An accident is not inevitable in either case. The aircraft may still be on the ground or may be able to glide to a safe landing. The chemicals may be released at a time when no wind is present to blow them into a populated area, and they may simply dissipate into the atmosphere. In neither case has any loss occurred.[2]

A loss results from the combination of a hazardous system state and environmental state:

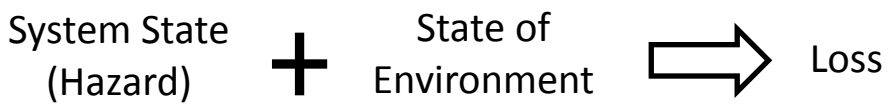

The introduction of the term "hazard" is important in engineering. The engineers or system designers and the system operators only have under their control the system itself and not the environment. Because the goal is to prevent hazards, that goal is achievable only if the occurrence of the hazard is under someone's control. When designing the chemical plant, the designer does not have any control over which way the wind is blowing when chemicals are released into the environment. The only thing they and the operators can do is to try to prevent the release itself through the design or operation of the system, in other words, by controlling the hazard or system state. An air traffic control system can control whether an aircraft enters a region with potentially dangerous weather conditions, but air traffic control has no control over whether the aircraft is hit by lightning if it does enter the region. The aircraft designers have control over whether protection against lightning strikes is included in the aircraft

[2] One might argue that chemicals have been wasted but that would have to be included in the definition of a loss for the chemical plant and thus the hazard would be the chemical plant being in a state where chemicals could be released and wasted.

design, but not whether the aircraft will be struck by lightning. Therefore, when identifying system hazards, think about what things are under our control that could, in some particular environmental conditions, potentially lead to an accident. If no such environmental conditions are possible, then there is no hazard.[3]

Now let's go on to something that is hopefully more interesting.

[3] Some fields define "hazard" differently than system engineering. For example, in aviation, a mountain may be called a hazard because an airplane can be flown into it. But the goal in engineering is to eliminate or control a hazard. The mountain cannot, in most cases, be eliminated. The only thing the aircraft designers and operators have control over is staying clear of the mountain. Therefore, the hazard would be defined as violating minimum separation standards with dangerous terrain.

Chapter 3: Why Aren't We Learning Enough from Accidents and Incidents?

"A learning experience is one of those things that says, 'You know that thing you just did? Don't do that.'"

Douglas Adams, *The Salmon of Doubt*, William Heinemann Ltd, 2001.

While there are many limitations in the way we usually do accident causal analysis and learn from events, five may be the most important: root cause seduction and oversimplification of causal explanations, hindsight bias, superficial treatment of human error, a focus on blame, and the use of models of causality that do not fit today's world.

Root Cause Seduction and Oversimplification of Causality

Humans appear to have a psychological need to find a straightforward and single cause for a loss, or at least a very limited number of causes. John Carroll calls this "Root Cause Seduction." We want simple answers to complex problems. Not only does that make it easier to devise a response to a loss, but it provides a sense of control. If we can identify one cause or even a few that are easy to fix, then we can "solve the problem" and shift our attention to other concerns.

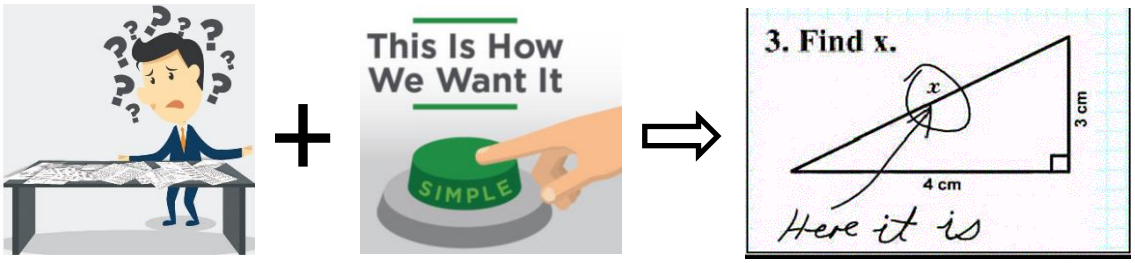

Figure 1: Root Cause Seduction leads nowhere.

The result of searching for a root cause and claiming success is that the problem is not fixed and further accidents occur. We end up in continual fire-fighting mode: fixing the symptoms of problems but not tacking the systemic causes and processes that allow those symptoms to occur. Too often we play a sophisticated "whack-a-mole" game and do not understand why the losses continue to occur. Enormous resources may be expended with little return on the investment.

Figure 2: Playing Whack-a-Mole

Here are some examples of oversimplification of causal analysis leading to unnecessary accidents. In the crash of an American Airlines DC-10 at Chicago O'Hare Airport in 1979, the U.S. National Transportation Safety Board (NTSB) blamed only a "maintenance-induced crack" and not also a design error that allowed the slats to retract if the wing was punctured. Because of this omission, McDonnell Douglas was not required to change the design, leading to future accidents related to the same design error.

In the explosion of a chemical plant in Flixborough, Great Britain, in June 1974, a temporary pipe was used to replace a reactor that had been removed to repair a crack. The crack itself was the result of a poorly considered process modification. The bypass pipe was not properly designed (the only drawing was a sketch on the workshop floor) and was not properly supported (it rested on scaffolding). The jury-rigged bypass pipe broke, and the resulting explosion killed 28 people and destroyed the site. The accident investigators devoted much of their effort to determining which of two pipes was the first to rupture. The British Court of Inquiry concluded that "The disaster was caused by a coincidence of a number of unlikely errors in the design and installation of a modification" (the bypass pipe) and that "such a combination of errors is very unlikely ever to be repeated." [245].

Clearly, however, the pipe rupture was only a small part of the cause of this accident. A full explanation and prevention of future such losses required an understanding, for example, of the management practices of running the Flixborough plant without a qualified engineer on site and allowing unqualified personnel to make important engineering modifications without properly evaluating their safety, as well as storing large quantities of dangerous chemicals close to potentially hazardous areas of the plant and so on. The British Court of Inquiry investigating the accident amazingly concluded that "there were undoubtedly certain shortcomings in the day-to-day operations of safety procedures, but none had the least bearing on the disaster or its consequences and we do not take time with them." Fortunately, others ignored this overly narrow view, and Flixborough led to major changes in the way hazardous facilities were allowed to operate in Britain.

In many cases, the whack-a-mole approach leads to so many incidents occurring that they cannot all be investigated in depth, and only superficial analysis of a few are attempted. If instead, a few were investigated in depth and the systemic factors fixed, the number of incidents would decrease by orders of magnitude.

In some industries, a conclusion is reached when accidents keep happening that accidents are inevitable and that providing resources to prevent them is not a good investment. Like Sisyphus, they feel like they are rolling a large boulder up a hill with it inevitably crashing down to the bottom again until they finally give up, decide that their industry is just more dangerous than the others that have better accident statistics, and conclude that accidents are the price of productivity. Like those caught in any vicious circle, the solution lies in breaking the cycle, in this case by eliminating oversimplification of causal explanations and expanding the search for answers beyond looking for a few root causes.

Accidents are always complex and multifactorial. Almost always there is some physical failure or physical equipment that had flaws in its design, operators who at the least did not prevent the loss or whose behavior may have contributed to the hazardous state, flawed management decision making, inadequate engineering development processes, safety culture problems, regulatory deficiencies, etc. Jerome Lederer, considered the Father of Aviation Safety, wrote:

> ”Systems safety covers the total spectrum of risk management. It goes beyond the hardware and associated procedures of systems safety engineering. It involves: attitudes and motivation of designers and production people, employee/management rapport, the relation of industrial associations among themselves and with government, human factors in supervision and quality control, documentation on the interfaces of industrial and public safety with design and operations, the interest and attitude of top management, the effects of the legal system on accident investigations and exchange of information, the certification of critical workers, political considerations, resources, public sentiment and many other non-technical but vital influences on the attainment of an acceptable level of risk control. These non-technical aspects of system safety cannot be ignored.”[4]

Our accident investigations need to potentially include all of these factors and more. This handbook will show you how.

Hindsight Bias

A lot has been written about the concept of hindsight bias. At the risk of oversimplifying, hindsight bias means that after we know that an accident occurred and have some idea of why, it is psychologically impossible for people to understand how someone might not have predicted the events beforehand. After the fact, humans understand the causal connections and everything seems obvious. We have great difficulty in placing ourselves in the minds of those involved who have not had the benefit of seeing the consequences of their actions (see Figure 3).

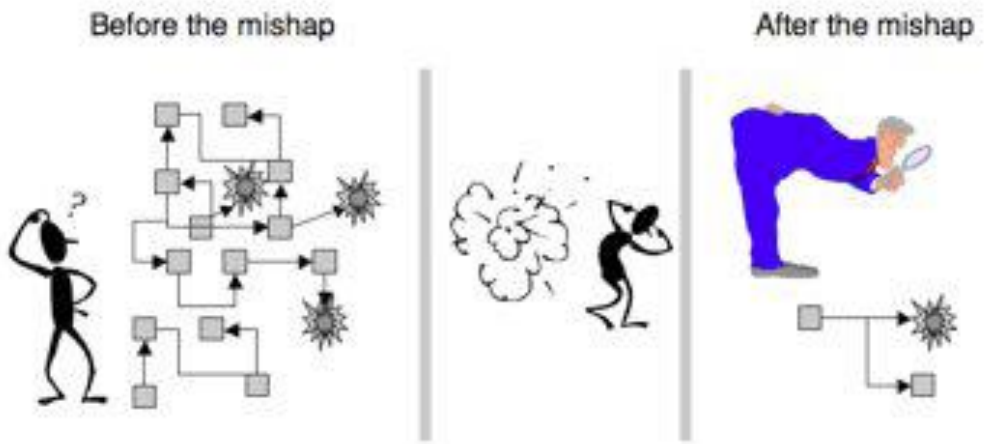

Figure 3: A graphical depiction of hindsight bias.
[Figure attributable to Richard Cook or Sidney Dekker]

Hindsight bias is usually found throughout accident reports. A glaring clue that hindsight bias is involved is when you see the words “he/she should have…,” “he/she could have…,” or “if only he/she

[4] Jerome Lederer, How far Have we come? A look back at the leading edge of system safety eighteen years ago, *Hazard Prevention*, page 8, May/June 1986.

would have ….” Here are a couple of examples, one in a chemical plant and the other involving an aircraft.

After an accident involving the overflow of SO^2 (sulfur dioxide) in a chemical plant, the investigation report concluded that “*The Board Operator should have noticed the rising fluid levels in the tank*.” [emphasis added] Sounds bad, right? Let’s examine that conclusion.

The operator had turned off the control valve allowing fluid to flow into the tank, and a light came on saying it was closed. All the other clues that the operator had in the control room showed that the valve had closed, including the flow meter, which showed that no fluid was flowing. The high-level alarm in the tank did not sound because it had been broken for 18 months and was never fixed. There was no indication in the report about whether the operators knew that the alarm was not operational. Another alarm that was supposed to detect the presence of SO^2 in the air also did not sound until later.

One alarm did sound, but the operators did not trust it as it had been going off spuriously about once a month and had never in the past signaled anything that was actually a problem. They thought the alarm resulted simply from the liquid in the tank tickling the sensor. While the operators could have used a special tool in the process control system to investigate fluid levels over time (and thus determine that they were rising), it would have required a special effort to go to a page in the automated system to use the non-standard tool. There was no reason to do so (it was not standard practice) and there were, at the time, no clues that there was a problem. At the same time, an alarm that was potentially very serious went off in another part of the plant, which the operators investigated instead. As a result, the operators were identified in the accident report as the primary cause of the SO^2 release.

It’s interesting that the report writers could not, even after careful study after the release, explain why the valve did not close and the flow meter showed no flow; in other words, why the tank was filling when it should not have been. But the operators were expected to have known this without any visible clues at the time and with competing demands on their attention. This is a classic example of the investigators succumbing to hindsight bias. The report writers knew, after the fact, that SO^2 had been released and assumed the operators should have somehow known too.

The words “should have” or their equivalent may not appear in the report, but hindsight bias may still be at work. As an example, one of the four probable causes cited in the accident report of the American Airlines B757 crash while approaching Cali, Columbia in 1995 was “*Failure of the flight crew to discontinue the approach into Cali, despite numerous cues alerting them of the inadvisability of continuing the approach*.” In fact, the “cues” were only cues in hindsight if one already knows that a crash has occurred.

In summary, hindsight bias occurs because, after an accident, it is easy to see where people went wrong and what they should have done or avoided doing. It is also easy to judge about missing a piece of information that turns out to be critical only after the causal connections for the accident are made. It is almost impossible to go back and understand how the world looked to somebody not having knowledge of the later outcome.

How can hindsight bias be avoided, given that it is a natural result of ‘after the fact” reasoning? It takes some effort and a change in the way we think about causality. Instead of spending our time focused on identifying what people did wrong when analyzing the cause of an accident, we instead need to start from the premise that the operators were not purposely trying to cause a loss but instead were trying to do the right thing. Learning can occur when we focus on identifying not what people did wrong

but *why it made sense to them at the time to do what they did*.[5] CAST requires answering this type of question and leads to identifying more useful ways to prevent such behavior in the future.

Unrealistic Views of Human Error

A treatise on human factors is not appropriate here. Many such books exist. But most accident analyses start from a belief that operator error is the cause of most incidents and accidents.[6] Therefore, it follows that the investigation should focus primarily on the operator. An assumption is made that the operator must be the cause and, unsurprisingly, the operator is then the focus of attention in the accident analysis and identified as the cause. Once the operator is implicated, the recommendations emphasize doing something about the operator (punish them, fire them, retrain the particular operator or all operators not to do the same thing again). The emphasis on human error as the cause of accidents is partly due to the "bad apple" theory, which arose a hundred years ago and was thoroughly discredited scientifically about seventy years ago. Unfortunately, it still persists. Appendix C provides more information about it. Heinrich also promulgated this theory around the same time.

Alternatively, or in addition, something may be done about operators in general. Their work may be more constrained by creating more rules and procedures—which may be impossible or unrealistic to expect them to always follow or which may themselves lead to an accident. Or the response may be to marginalize the operators by adding more automation. Adding more automation may introduce more types of errors by moving operators farther from the process they are controlling. Most important, by focusing on the operators, the accident investigation may ignore or downplay the systemic factors that led to the operator behavior and the accident.

As just one example, many accident investigations find that operators had prior knowledge of similar previous occurrences of the events but never reported them in the incident reporting system. In many cases, the operators did report them to the engineers who they thought would fix the problem, but the operators did not use the official incident-reporting system. A conclusion of the report then is that a cause of the accident was the operators not using the incident-reporting system, which leads to a recommendation to make new rules to enforce that operators always use it and perhaps recommend providing additional training in its use.

In most of these cases, however, there is no investigation of *why* the operators did not use the official reporting system. Often their behavior results from the system being hard to use, including requiring the operators to find a seldom-used and hard to locate website with a clunky interface. Reporting events in this way may take a lot of time. The operators never see any results or hear anything back and assume the reports are going into a black hole. It is not surprising then that they instead report the problem to people who they think can and will do something about it. Fixing the problems with the design of the reporting system will be much easier and more effective than simply emphasizing to operators that they have to use it.

A system's view of human error starts from the assumption that all behavior is affected by the context (system) in which it occurs. Therefore, the best way to change human behavior is to change the system in which it occurs. That involves examining the design of the equipment that the operator is using, carefully analyzing the usefulness and appropriateness of the procedures that operators are given

---

[5] For more about this, see Sidney Dekker, *The Field Guide to Understanding Human Error*, Ashgate Publishers, 2002.

[6] Much research is published that concludes that operators are the cause of 70-90% of accidents. The problem is that this research derives from looking at accident reports. Do the conclusions arise from the fact that operators actually are the primary cause of accidents or rather that they are usually blamed in the accident reports? Most likely, the latter is true. At best, such conclusions are not justified by simply looking at accident reports.

to follow, identifying any goal conflicts and production pressures, evaluating the impact of the safety culture in the organization on the behavior, and so on.

Violating safety rules or procedures is interesting as it is commonly considered prima facie evidence of operator error as the cause of an accident. The investigation rarely goes into why the rules were violated. In fact, rules and procedures put operators and workers into an untenable situation where they are faced with a "following procedures" dilemma. Consider the model in Figure 4.

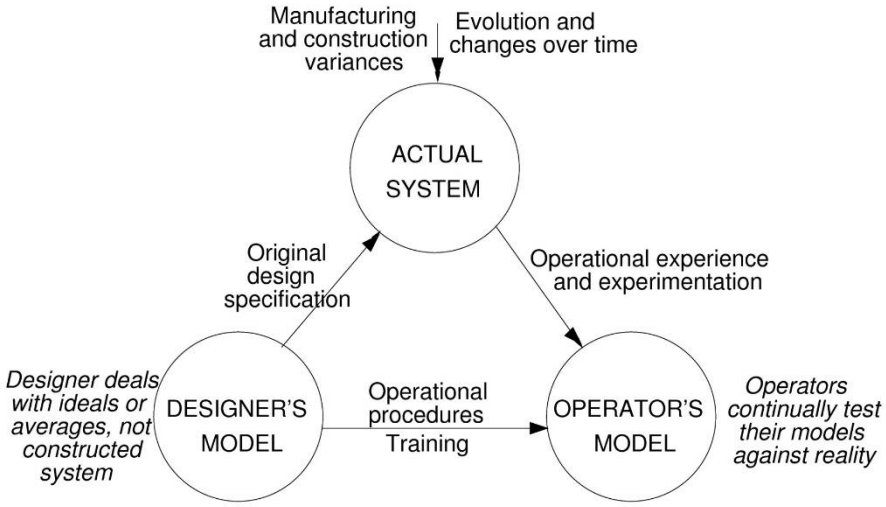

Figure 4. The Following Procedures Dilemma

There are always at least two mental models of the actual system: the designer's model and the operator's model. The designer provides or complies with the original design specification and the operational procedures and training guidance. The designer deals with ideals or averages (the ideal material or the average material) and assumes that the actual system will start out satisfying the original design specification and remain that way over time. The operational procedures and training are based on that assumption. In reality, however, there may be manufacturing and construction variances during the initial construction. In addition, the system will evolve and its environment will change over time.

The operator, in contrast, must deal with the actual system as it exists at any point in time, not the system that was originally in the designers' minds or in the original specifications. How do operators know what is the current state of the system? They use feedback and operational experience to determine this state and to uncover mistaken assumptions by designers. Often, operators will test their own mental models of the actual system state by performing "experiments" or informal tests. They are continually testing their own models of the system behavior and current state against reality.

The procedures provided to the operators by the system designers may not apply when the system behaves differently than the operators (and designers) expected. For example, the operators at Three Mile Island recognized that the plant was not behaving the way they expected it to behave. They could either continue to follow the utility-provided procedures or strike out on their own. They chose to follow the procedures, which after the fact were found to be wrong. The operators received much of the blame for the incident due to them following those procedures. In general, operators must choose between:

1. Sticking to procedures rigidly when cues suggest they should instead be adapted or modified, or
2. Adapting or altering procedures in the face of unanticipated conditions.

The first choice, following the procedures they were trained to follow, may lead to unsafe outcomes if the trained procedures are wrong for the situation at hand. They will be blamed for their inflexibility and applying rules without understanding the current state of the system and conditions that may not have been anticipated by the designers of the procedures. If they make the second choice, adapting or altering procedures, they may take actions that lead to accidents or incidents if they do not have complete knowledge of the current circumstances and system state and what might happen if they violate the written procedures. They will then be blamed for deviations and rule violations.

Following rules and procedures does not necessarily guarantee safety. Instead, safety comes from people being skilled in judging when and how to apply them. The traditional approach to safety contends that safety improvements come from organizations telling people to follow procedures, enforcing them, and assigning blame to people when accidents occur and there was a rule or procedure violation. In contrast, a system's approach to safety assumes that safety improvements come from organizations monitoring and understanding the gap between written procedures and practice (behavior) and updating procedures and rules accordingly. After accidents or incidents where written procedures or rules are violated, the focus should be placed on determining why the operators felt the need to disregard the procedures or rules. Simply concluding that the operators were at fault because they did not follow the written rules provides no useful information.

In general, the role of operators in modern systems is changing. Rather than directly controlling the system, humans are increasingly supervising automation, which is in fact implementing most of the detailed control tasks. At the same time, software is allowing enormously complex systems to be created, which are stretching the ability of people to understand them and leading to human behavior that under some conditions could be unsafe. In addition, systems are sometimes designed without using good human-centered and human-factors design principles. The result is that we are designing systems in which operator error is inevitable and then blaming accidents on operator error rather than designer error.

Human error is a symptom, not a cause. Consider another one of the four probable causes cited in the accident report for the American Airlines Cali accident: "Failure of the flight crew to revert to basic radio navigation at the time when the FMS-assisted navigation became confusing and demanded an excessive workload in a critical phase of flight." Notice that it is deemed to be the flight crew's fault for being confused and also for choosing to use poorly designed software (the FMS or Flight Management System) and not the fault of the software for being confusing and demanding an excessive workload. When accident reports state causes in this way, learning is inhibited and attention is directed to the wrong aspects of the causal scenario.

A systems approach to accident causation starts from the premise that human error is a symptom of a system that needs to be redesigned. Accident analysis should identify the design flaws and recommend ways they can be fixed, not blame the operators for the consequences of those design flaws.

Blame is the Enemy of Safety

Blame is a legal or moral concept, not an engineering one

Focusing on blame seriously hinders what we learn from accidents. The goal of courts is to establish blame and liability. The goal of engineering, in contrast, is to understand *why* accidents occur so they can be prevented, not to establish blame and decide who was responsible.

A focus on blame in accident analysis has many unfortunate aspects that reduce learning from accidents and impedes preventing future ones. One result is that important information is often hidden: Those involved resort to pointing fingers at everyone else and searching for someone else to blame. The

search for causes devolves to identifying the immediate actors in the event chain, usually the human operators or low-level managers, who obviously participated in the events and have no way to deflect attention onto others. The spotlight then is placed on the aspects of the loss that are least likely to provide important information about preventing future accidents.

C.O. Miller and Gerry Bruggink distinguished between an accusatory approach to causal analysis and an explanatory one. Here is a short exercise for the reader of this handbook. Consider the following conclusion from an accident report:[7]

Exercise Part 1. Description of the cause of the accident:

WHO

Accident Board **A** determined the probable cause of this accident was:

1. The flight crew's failure to use engine anti-icing during ground operations and takeoff.
2. Their decision to take off with snow/ice on the airfoil surfaces of the aircraft, and
3. The captain's failure to reject the takeoff during the early stage when his attention was called to anomalous engine instrument readings.

WHY

Contributing Factors:

1. The prolonged ground delay between de-icing and receipt of ATC clearance during which the airplane was exposed to continual precipitation
2. The known inherent pitch-up characteristics of the B-737 aircraft when the leading edge is contaminated with even small amounts of snow or ice, and
3. The limited experience of the flight crew in jet transport winter operations.

Using this description of the cause of the accident, where would you assign responsibility for this loss? What recommendations would you provide as a result? Are there additional questions you might ask during the investigation?

Discussion of your answer:

Did you assign responsibility to the flight crew? Why? What was the focus of the recommendations that you generated? Did they involve flight crew training and procedures? What other causal factors and recommendations did you generate? You might have mentioned something about the ground delay and avoiding that, but most likely all of your causes and recommendations involved the flight crew. That would not be surprising as their actions are identified in the analysis as the probable cause of the accident.

[7] Adapted from C.O. Miller, "Down with 'Probable Cause'!" International Society of Air Safety Investigators Seminar, Canberra, Australia, November 1991.

Exercise Part 2: Now consider a second description of the cause for the same accident:

WHAT

Based on the available evidence, Accident Board **B** concluded that a thrust deficiency in both engines, in combination with contaminated wings, critically reduced the aircraft's takeoff performance, resulting in a collision with obstacles in the flight path shortly after liftoff.

WHY

Reasons for the thrust deficiency:

1. Engine anti-icing was not used during takeoff and was not required to be used based on the criteria for "wet snow" in the aircraft's operations manual.
2. The engine inlet probes became clogged with ice, resulting in false-high thrust readings.
3. One crew member became aware of anomalies in cockpit indications but did not associate these with engine inlet probe icing.
4. Despite previous incidents involving false thrust readings during winter operations, the regulator and the industry had not effectively addressed the consequences of blocked engine inlet probes.

Reasons for the wing contamination:

1. Deicing/anti-icing procedures.
2. The crew's use of techniques that were contrary to flight manual guidance and aggravated the contamination of the wings.
3. ATC procedures that resulted in a 49-minute delay between departure from the gate and takeoff clearance.

Now, to whom or what would you assign responsibility for this accident? Was your list larger and different than the list you generated in exercise 1? Are there additional questions that you might like to investigate? Do you now think the cause was more than just "flight crew failures"? Did you identify other factors? What recommendations would you create from this analysis? How do the recommendations you generated differ from those for Exercise 1? Why? Did you generate design recommendations or recommendations involving groups or people other than the flight crew?

The first description of the accident (in Exercise 1) was *accusatory*. It focuses on the operators and their role in the loss, that is, on who was responsible. The second description (in Exercise 2) was *explanatory*. It does not focus on who but instead on what and why. Very different recommendations will result from each of these, although they are describing the same accident. We will learn more and find more causal factors if we use an explanatory approach rather than an accusatory one, that is, if we focus on what and why and not who.

Accusatory:		**Explanatory**:
Who	**vs.**	What
Why		Why

Figure 5: Two opposing views of accident explanation

Another thing to notice in these two exercises is the use of the word "failure" in the first but not in the second. Failure is a pejorative word, that is, it involves judgement and assignment of blame. Consider the two following two statements, differing only in the use of the word "failure":

1. The captain's failure to reject the takeoff during the early stage when his attention was called to anomalous engine instrument readings.
2. The captain did not reject the takeoff during the early stage when his attention was diverted to anomalous engine instrument readings.

The first suggests a conclusion that the captain did something wrong. Further exploration is not encouraged because it is stated as a conclusion. A judgment has been made about the captain's actions. A cause has been determined, and that cause was a failure on the part of the captain. The second is a simple statement of what the captain did without any judgment in it. It encourages further exploration about why the captain's attention might have been drawn to the anomalous engine instrument readings and, even more important, why anomalous readings were produced.

Is this really a problem in accident reports? In fact, the conclusion of "failure" pervades most accident reports. Here is an example of the conclusions from an aircraft CFIT (Controlled Flight into Terrain) accident at Birmingham Shuttlesworth International Airport,[8] which is typical. The NTSB concluded that:

> "The probable cause of this accident was the flight crew's continuation of an unstabilized approach and their failure to monitor the aircraft's altitude during the approach, which led to an inadvertent descent below the minimum approach altitude and subsequently into terrain.

The report also concludes that contributing to the accident were:

(1) the flight crew's failure to properly configure and verify the flight management computer for the profile approach;
(2) the captain's failure to communicate his intentions to the first officer once it became apparent the vertical profile was not captured;
(3) the flight crew's expectation that they would break out of the clouds at 1,000 feet above ground level due to incomplete weather information;
(4) the first officer's failure to make the required minimums callouts;
(5) the captain's performance deficiencies likely due to factors including, but not limited to, fatigue, distraction, or confusion, consistent with performance deficiencies exhibited during training; and
(6) the first officer's fatigue due to acute sleep loss resulting from her ineffective off-duty time management and circadian factors.

[emphasis added]

Notice that the conclusions about the probable cause and contributing factors for this accident identify only flight crew behavior and the events that reflect flight crew "failures." *Why* the flight crew behaved the way they did is not included (except for fatigue, which does not fully explain it). One contributory factor does mention a reason, i.e., (3) the flight crew's expectation that they would break out of the clouds at 1,000 feet above ground level due to incomplete weather information. But the emphasis here is on the flight crew behavior and expectations and not why incomplete weather information was provided to the crew. The system design flaws related to providing weather and needing to be fixed are not mentioned in the accident causal summary.

[8] National Transportation Safety Board, Crash During a Nighttime Nonprecision Instrument Landing, UPS Flight 1354, Birmingham, Alabama, August 14, 2013, Accident Report NTSB/AAR-14/02, 2014.

Alternatively, a systems approach using CAST would not only look at what the pilots did that contributed to the accident but, more important, *why* they believed it was the right thing to do at that time.[9] The official conclusions omit most of the information that would be useful in preventing such accidents in the future. We don't learn why the flight crew behaved the way they did because the explanation stops with the conclusion that they "failed." In addition, the entire system for preventing CFIT (which has been a focus of accident prevention for many years) needs to be examined—not just the pilot behavior. Some of the many contributors to this accident were the design of the automation, landing on a runway without an instrument landing system, lack of an air traffic control alert to indicate that the aircraft was too low, maintenance practices at the airport, the ground proximity warning system not providing an alert until too late (due in part to the airline not upgrading the warning system with the new software provided), and a dozen other factors that are more likely to lead to useful recommendations than simply identifying the "failures" of the flight crew.

The use of the word "failure" pervades the causal description of most accident reports. It is erroneously applied not only to humans but to software, operators, and management decision making.

Software does not "fail"; it simply executes the logic that was written. There needs to be an examination of why unsafe software was created—usually it can be traced to requirements flaws—and recommendations involving improvement of the process that produced the unsafe software. Concluding that the "software failed" makes no technical sense and provides no useful information.

Humans also do not "fail" (unless their heart stops). They simply react to the situations in which they find themselves. What they did may, in hindsight, turn out to have been the wrong thing to do. But why it seemed to them to be the right thing at the time needs to examined in order to make useful recommendations. Again, simply recounting what people did wrong provides no useful information beyond creating a convenient scapegoat for the accident. Understanding why they behaved the way they did will focus the accident analysis in a useful direction.

And finally, companies do not fail unless they go out of business. A typical example found in accident reports is a statement that "Company X failed to learn from prior events." Companies are pieces of paper describing legal entities. Documents do not learn or, for that matter, fail. The company may be made up of hundreds, thousands, or even tens of thousands of employees. More useful would be to determine why learning did not take place and to ask questions such as: Is there a safety information system that captures prior adverse events so that learning can occur? If so, were the prior events recorded in the safety information system? If not, then why not? Were they just missed accidentally, was there no process to include them, or was the process not followed? If the latter, why was the process not followed? If the previous events <u>were</u> recorded, then why were they not used? Were they hard to find or was enough information not included to prevent repetition? Were there procedures in place to require using recorded information? And so on. Answering these questions will provide information useful for creating recommendations that will improve future learning from events and prevent repetition of the losses. Simply concluding that "the company failed to learn from events" only assigns blame without providing the information necessary to improve the situation in the future.

Eliminating the use of "failure" will greatly enhance learning from accidents. By not short-circuiting the search for understanding by concluding there was a "failure," we could immeasurably improve accident investigation and learning and, hopefully, safety.

More generally, to increase learning from losses, we need to eliminate the focus on assigning blame in accident reports. This may be hard to do: there seems to be a human need to find someone responsible for a tragedy. But a choice has to be made between personal satisfaction in identifying a villain or seizing the opportunity to prevent more such accidents in the future. Engineers should leave

[9] Sidney Dekker, *A Field Guide to Understanding Human Error*, London: Ashgate, 2006.

blame to the legal system. And companies and others should not be allowed to use the official accident investigations to shift liability away from themselves and onto someone else. A report that simply describes what happened, including all the contributions to the events and why, in a blame-free manner will shift the effort to assign or avoid liability to the legal process where it belongs.

Use of Inappropriate Accident Causality Models

One other factor has a major impact on what and how much we learn from accident investigations. Many of the problems described in the previous sections stem from the definition of "cause" used in the accident analysis. It is therefore worth considering this definition a bit more. The definition of "cause" depends on the accident causality model being used in the investigation. This section answers the questions:

- *What is an accident causality model?*
- *What causality model is almost always used today?*
- *What are the limitations of the traditional causality model?*
- *Is there an alternative?*

What is an accident causality model?

A causality model simply explains why something happened. For accidents, the model is used to explain why the accident occurred. We all have a model in our heads of how and why we think accidents occur. We may not consciously think about it, but it influences our thinking. When we try to learn from and prevent accidents, we use that model in the process. If the model limits the causal mechanisms considered, then it will reduce learning.

What causality model is almost always used today?

The most common model of causality assumes that accidents are the result of a chain of events, where the events usually involve failures of system components and each failure or event is the direct cause of the one preceding it.

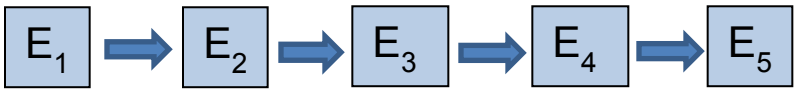

This model has been given various names and analogies used for depictions of the events (Figure 6 and Figure 7), but the underlying model is the same, that is, a chain of failure events.

Figure 6: Heinrich's Domino Model, 1932.

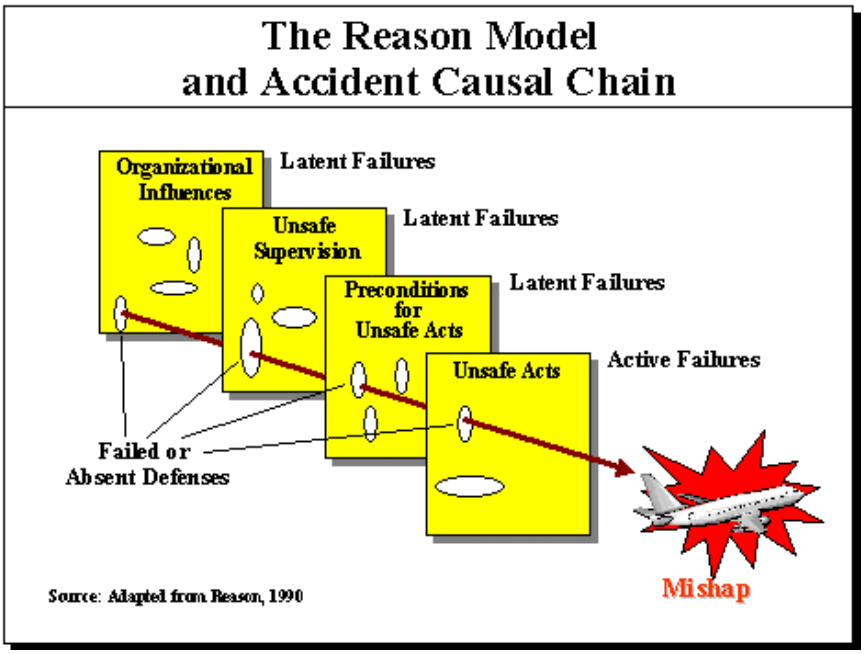

Figure 7: Reason's Swiss Cheese Model

As an example, consider the Bhopal (India) chemical plant release of methyl isocyanate (MIC), which killed and seriously injured tens of thousands of people. The chain of events might be stated as the following:

E1: A worker washes out pipes without inserting a slip blind to keep water isolated from MIC.

E2: Water from the pipe washing leaks into the MIC tank (MIC is explosive in the presence of water.) and pressure rises in the tank.

E3: The control room operator does not open the valve to a secondary relief tank to relieve pressure because the gauges are out of order. He thinks the relief tank is full.

E4: An explosion occurs.

E5: The automated relief valve opens (as designed) in the presence of high pressure.

E6: None of the protection devices (flare tower, vent scrubber, and water curtain) prevent the release of the MIC and it is vented into the air.

E7: Wind carries MIC into the highly populated area around the plant.

From this chain of events, one event is usually selected as the root cause. In this case, both Union Carbide (the owner of the plant) and the Indian subsidiary running the plant pointed at the low-level worker who washed out the pipes without inserting a slip blind, and he was actually put in jail. Note, however, that this choice is arbitrary. Any of the other events could have been selected, such as the operators not opening the valve to the relief tank or the protection devices not working.

The chain of events here could have been followed farther back, but the chaining usually stops when a convenient person or event is found upon which to place the blame, in this case the pipe washer. The selected root cause event usually involves an operator action or inaction that is close to the time of the actual loss. Putting the worker in jail may have satisfied the goals of most everyone else involved (who were quite happy to deflect the blame away from themselves) but does not help to identify and eliminate the true causes of the accident, which were much more complex. Basically, to learn from accidents and prevent future ones, we need to know not just what events occurred but why. And the assumption that the events always occur in a simple causal chain is untrue.

For example, why did the pipe washer not insert the slip blind? In fact, the worker who had been assigned the task of washing out 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 slip blinds 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 operator, but that position had been eliminated in a cost-cutting effort. Using the chain of events model, we could, of course, simply insert a previous event, such as firing the second shift supervisor, as the root cause. But that doesn't solve the problem because the chain still does not explain why this accident occurred. The accident, as is the case for most accidents when they are carefully investigated, turned out to be much more complex than just a simple chain of events. A major limitation in understanding this complex set of events and avoiding future such events derives from the oversimplification implicit in using a chain of events at all. Given the design and operating conditions of the Bhopal plant, in fact, an accident was waiting to happen:

> "However [water] got in, it would not have caused the severe explosion had the refrigeration unit [required to keep the MIC at a non-explosive temperature] not been disconnected and drained of freon [to save money], 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."[10]

And this is just the beginning. Among a large number of other things, a safety audit two years before had noted many problems at the plant, including all the ones implicated in the loss, but the identified deficiencies were never fixed. The details are unimportant here. The basic problem is that causality in complex losses cannot be captured by simple chains of events.

[10] Charles Perrow, The Habit of Courting Disaster, *The Nation*, October, 1986, p. 349

What are the limitations of this traditional causality model?

One of the most important limitations is an assumption about the linearity of accident causes. Note also that aside from the influence of one event on the next event in the chain, there is an assumption that the events are independent. In fact, there are almost always systemic causes that impact all the events in the chain.

Consider again the Bhopal accident, whose cause was identified as the worker who washed out the pipes without inserting protection against a leaky valve. Many of the poor conditions at the plant were allowed to persist because of financial pressures: Demand for MIC had dropped sharply, leading to reductions in production and pressure on the company to cut costs. 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 showed up for his shift and thus he was not replaced, and nobody inserted the protection device.

As the Bhopal plant lost money, many of the skilled workers left for more secure jobs. They either were not replaced or were replaced by unskilled workers (like the pipe washer who was blamed for the accident). Staff educational standards were reduced along with training. Management and labor problems followed the financial losses, leading to low morale at the plant.

These are only a few of the factors involved in this catastrophe, which also included other technical and human errors in 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 such a dangerous chemical was being made 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. When all the factors are considered, the pipe washing operator looks less and less likely to be the "root" or "probable" cause of this accident. Do you think that putting the pipe washer in jail (as occurred after the accident) would be effective in preventing future accidents in this plant given all the problems stated so far (and there were many more)?

Implicit in the chain of events model is an assumption of direct causality between the events in the chain, that is, A→B, which states that A is a necessary and sufficient cause of B. The contrapositive of this logical statement implies that if B does not occur, then A has not occurred. While this makes for elegant formal logic, it inhibits the comprehensive identification of accident causes when applied to the messy real world in which we all live. For example, using the contrapositive and formal definitions of event chains, it is not possible to make the statement "Smoking causes lung cancer," because not everyone who smokes gets lung cancer and not everyone who gets lung cancer smokes. In fact, the Tobacco Lobby used the contrapositive argument for decades to prevent limitations being placed on smoking and cigarettes. We know that there is a connection, but it is clearly not a simple and direct one.

Many of the systemic causes described above are only indirectly related to the proximate events preceding the loss, which by using the contrapositive rule means they could not considered to be a cause. In fact, many of these factors had an impact on *all* the events in the chain, also negating the assumption that the events are unrelated except directly through a simple chain-like relationship. Identifying the chain of events provides little help in understanding why an accident occurred and how to prevent it in the future. At best, it is only a starting point for a causal analysis. Even then, it might limit the factors considered and bias the results.

Effectively preventing accidents in complex systems requires using accident models that include the social system as well as the technology and its underlying science. Without understanding the purpose, goals, and decision criteria used to construct and operate our systems, it is not possible to completely understand and most effectively prevent accidents.

Is there an alternative?

For a few simple accidents, the chain-of-events model may be adequate. To maximize learning from and preventing accidents in today's complex socio-technical systems, however, we need something more that will help us to more fully understand why the events occurred.

A new model of causality, which underlies CAST, provides that opportunity. STAMP (System-Theoretic Accident Model and Processes) is based on systems theory rather than traditional decomposition and reliability theory. In the latter, safety is essentially equated with the reliability of the system components.[11] Going into the technical details of STAMP is not necessary to use CAST. But it is useful to understand some important features of this new causality model.

CAST is rooted in systems thinking and systems theory. System theory, as used in engineering, was created after World War II to deal with the increased complexity of the systems starting to be built at that time. It was also independently created for biology in order to successfully understand the complexity of biological systems. In these systems, separation and independent analysis of interacting components (subsystems) distort the results for the system as a whole because the individual component behaviors are coupled in non-obvious ways. The first engineering uses of these new ideas were in the missile and early warning systems of the 1950s and 1960s. C.O. Miller introduced system theory concepts to safety starting in the 1950s. He also claimed to have been the first to use the name System Safety.

Sometimes "systems thinking" is used as an informal synonym for mathematically formal systems theory. Systems thinking is a holistic approach to analysis that focuses on how system components (physical devices, automation, humans, management, regulatory oversight, etc.) interrelate and behave over time. By using a systems approach to accident analysis, we get away from focusing on a few causal factors and instead provide a larger view of causality that can be used to identify powerful prevention measures.

Traditionally, complex systems have been decomposed into subcomponents with these subcomponents evaluated for some property of interest. The system may be separated into physical or functional properties or into chains of events over time. Then the results for each component are combined to get a result for the system as a whole. There are some assumptions upon which the utility of this process is based: the system components are assumed to operate independently and that separation into pieces (decomposition) does not distort the phenomenon being considered. This implies that the interactions among components can be examined pairwise, thus allowing them to be individually measured and combined to create a system value for the phenomenon. All of this also implies that the components (or events in a chain) are not subject to feedback loops and non-linear interactions.

System Theory was created in response to the fact that after WW II, our systems were becoming so complex that decomposition no longer worked (although that has not stopped people from using it). The assumptions underlying the decomposition approach just do not hold in today's software-intensive and complex systems. Instead, using System Theory:

- The system is treated as a whole, not as the sum of its parts. You have probably heard the common statement: "the whole is more than the sum of its parts."
- A primary concern is *emergent properties* (see Figure 8), which are properties that are not in the summation of the individual components but "emerge" when the components interact. Emergent properties can only be treated adequately by taking into account all their technical and social aspects. Safety is an emergent property as is security and other important engineered system properties.

[11] Nancy Leveson, *Engineering a Safer World*, MIT Press, 2012.

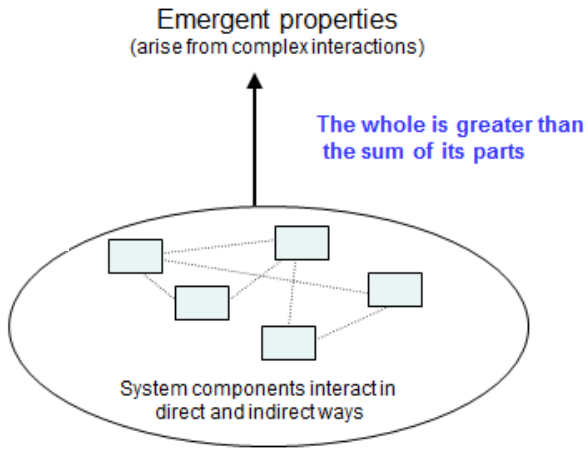

Figure 8: Emergent properties in system theory

- Emergent properties arise from the relationships among the parts of the system, that is, by how they interact and fit together.

If emergent properties arise from individual component behavior and complex interactions among components, then it makes sense that controlling emergent properties, such as safety, requires controlling both the behavior of the individual components and the interactions among the components. We can add a controller to the system to accomplish this goal. The controller provides control actions on the system and gets feedback to determine the impact of the control actions. In engineering, this is a standard feedback control loop (see Appendix E).

The controller *enforces constraints on the behavior of the system*. Example safety constraints might be that aircraft or automobiles must remain a minimum distance apart, pressure in deep water wells must be kept below a safe level, aircraft must maintain sufficient lift to remain airborne unless landing, toxic substances must never be released from a plant, and accidental detonation or launch of weapons must never occur.

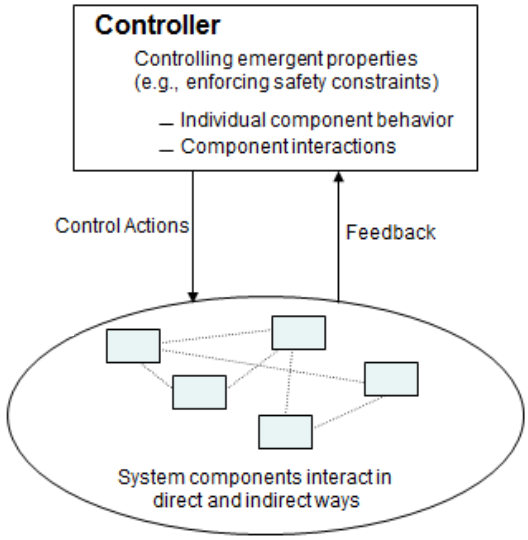

Figure 9: Controllers enforce constraints on behavior

Control is interpreted broadly and, therefore, includes everything that is currently done in safety engineering, plus more. For example, component failures and unsafe interactions may be controlled through design, such as using redundancy, interlocks, barriers, or fail-safe design. Safety may also be controlled through processes, such as development and training processes, manufacturing processes and procedures, maintenance processes, and general system operating processes. Finally, safety may be controlled using social controls including government regulation, culture, insurance, law and the courts, or individual self-interest. Human behavior can be partially controlled through the design of a societal or organizational incentive structure or other management processes.

The safety control structure enforces the safety constraints. Responsibility for enforcing specific behavior on the system is distributed throughout the system's control structure. Figure 10 is a generic example. Accidents result from inadequate enforcement of the system safety constraints in the operation of the system as a whole. In other words, weaknesses in the safety control structure leads to violation of system safety constraints and thus losses.

Handling complex sociotechnical systems requires modeling and analyzing both the social and technical aspects of the problem and allows a combined analysis of both. Figure 10 shows an example of a hierarchical safety control structure for a typical regulated industry in the United States. International controllers may be included. The operating process (the focus of most accident analysis) in the lower right of the figure makes up only a small part of the safety control structure.

There are two basic hierarchical control structures shown in Figure 10—one for system development (on the left) and one for system operation (on the right)—with interactions between them. Each level of the structure contains controllers with responsibility for controlling the interactions between and the behavior of the components on the level below. Higher level controllers may provide overall safety policy, standards, and procedures (downward arrows), and get feedback (upward arrows) about their effect through various types of reports, including incident and accident reports. For example, company management provides safety policy, standards, and resources and gets feedback in the form of reports about the operation of the company. The feedback provides the ability to learn and to improve the effectiveness of the safety controls.

Manufacturers must communicate to their customers the assumptions about the operational environment in which the original safety analysis was based, e.g., maintenance quality and procedures, as well as information about safe operating procedures. The operational environment, in turn, provides feedback to the manufacturer and potentially others, such as governmental authorities, about the performance of the system during operations. Each component in the hierarchical safety control structure has responsibilities for enforcing safety constraints appropriate for that component, and together these responsibilities should result in enforcement of the overall system safety constraints.

Note that the use of the term "control" does not imply a rigid command and control structure. Behavior is controlled not only by engineered systems and direct management intervention, but also indirectly by policies, procedures, shared value systems, and other aspects of the organizational culture. All behavior is influenced and at least partially "controlled" by the social and organizational context in which the behavior occurs. "Engineering" this context can be an effective way to create and change a safety culture, i.e., the subset of organizational or social culture that reflects the general attitude about safety by the participants in the organization or society.[12] More about the practical implications is provided later in the handbook.

[12] Edgar Shein, *Organizational Culture and Leadership*, San Francisco: Jossey Bass, 2004

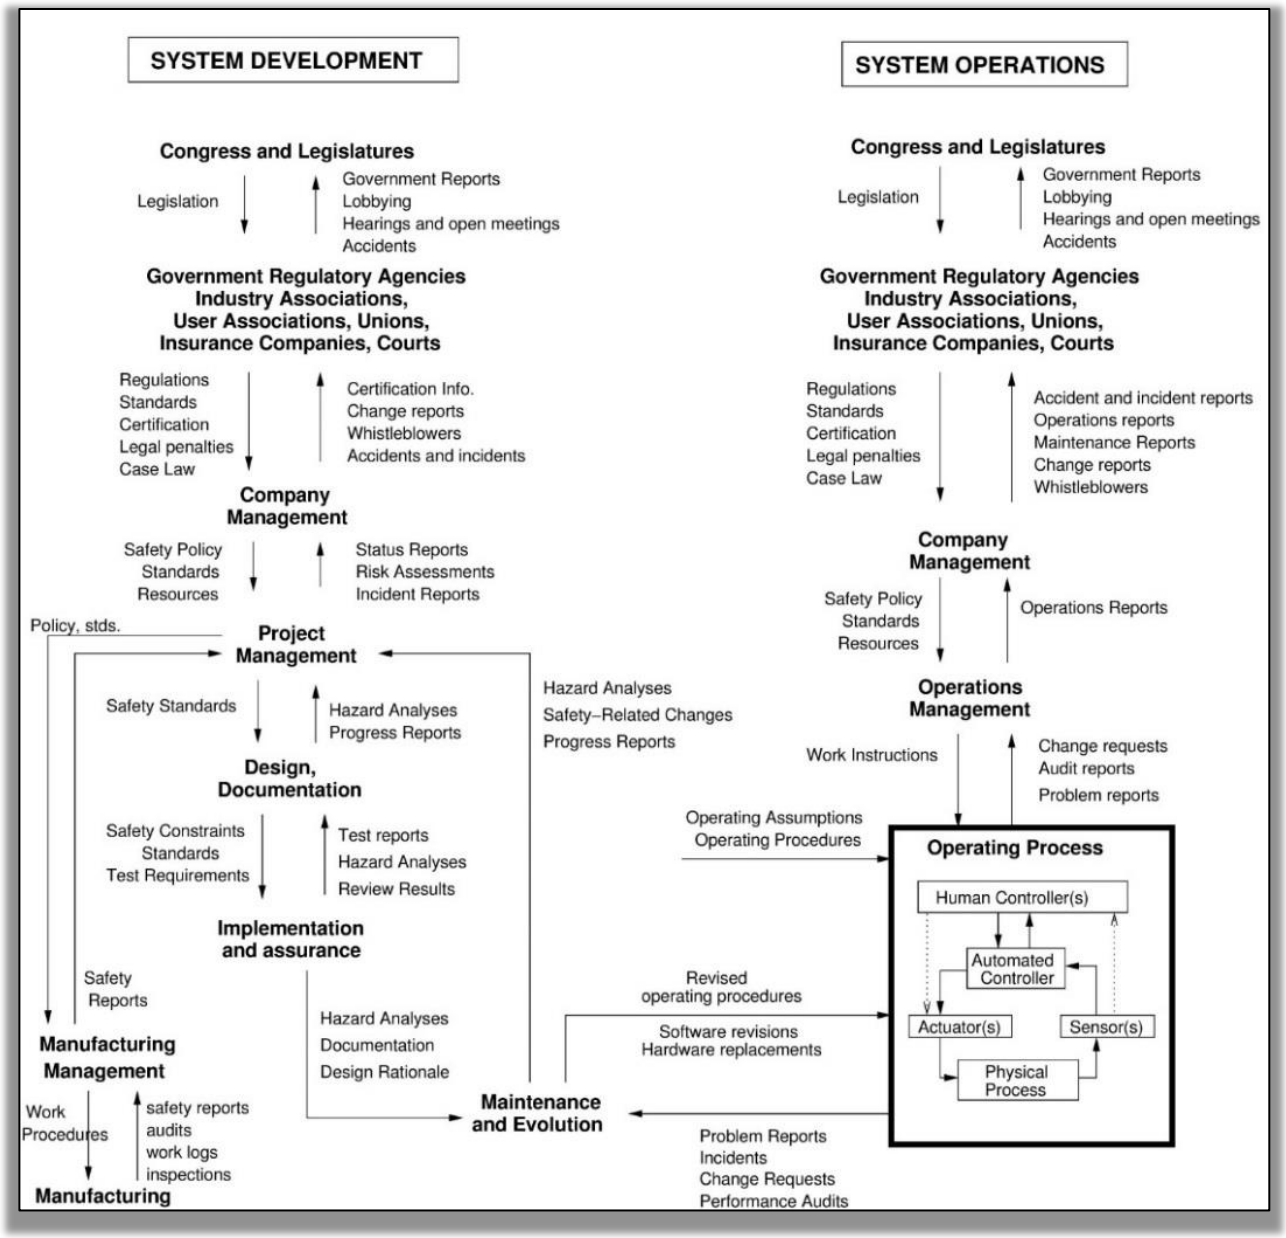

Figure 10: A generic safety control structure

As we have learned from major accidents, managerial and organizational factors and often governmental controls (or lack of them) are as important as technical factors in accident causation and prevention. For space reasons, Figure 10 emphasizes the high-level components of a safety control structure and not their detailed design. Each can be quite complex and needs to be examined comprehensively to understand why an accident occurred.

Figure 11 shows the basic form of the interactions between the levels of the control structure, where the controller imposes control actions on the controlled process. In engineering, this is called a feedback control loop. The standard requirements for effective management—assignment of responsibility, authority, and accountability—are part of the control structure design and specification as shown.

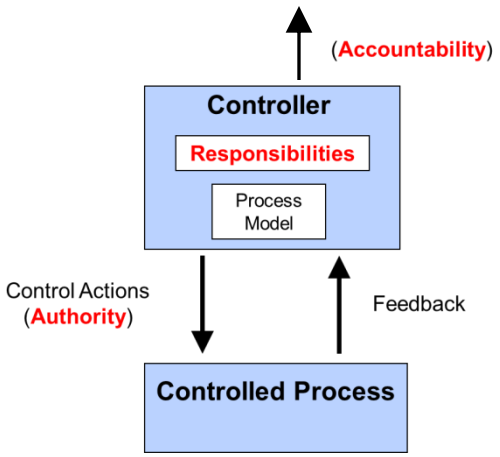

Figure 11: The basic building block for a safety control structure

The controller has responsibilities assigned to it with respect to enforcing the system safety constraints, for example, providing commands to an aircraft's control surfaces to control lift or, in a chemical plant, to open and close valves to maintain a safe level of pressure in a tank. The controller (which may be a human or automation) satisfies these responsibilities by issuing control actions on the process it is controlling (representing its authority). The controller determines what type of control actions are required to satisfy its responsibilities for preventing hazards given the current state of the controlled process, which is identified through feedback from the process being controlled.

As an example of a higher-level controller, the FAA has responsibilities related to overseeing the safety of flight in the U.S. The agency has various types of control actions to carry out its responsibilities, such as airworthiness advisory circulars and directives, FAA regulations, handbooks and manuals, Notices to Airmen (NOTAMs), policy and guidance, etc. Feedback comes in the form of reporting systems, accident and incident analysis, audits and inspections, etc. to determine the current state of safety of the air transportation system. Ultimately, the FAA is accountable to the Dept. of Transportation, Congress, and the public.

Feedback information is incorporated into the controller's model of the controlled process, called the *process model* or, if the controller is a human, it may be called the *mental model*. Accidents often result when the controller's process model becomes inconsistent with the actual state of the process and the controller provides unsafe control as a result. For example, the air traffic controller thinks that two aircraft are not on a collision course and does not change the course of one or both. Other examples are that the manager of an airline believing the pilots have adequate training and expertise to perform a particular maneuver safely when they do not, or a pilot thinking that de-icing has been accomplished when it has not.

In a simple feedback control loop, the controllers implement a control action, such as pulling back on the yoke in an aircraft or opening a valve in a plant. An "actuator" (not shown for simplicity) implements that control action in the controlled process. Feedback is usually designed into the system so that the controller can determine the effect of their control action and decide what to do next.

The importance of feedback becomes apparent here. Feedback informs the current process (or mental) model, which the controller uses to determine what control actions are needed. In engineering, every controller must contain a model of the controlled process in order to provide effective control. This process model or mental model includes assumptions about how the controlled process operates and the current state of the controlled process. It is used to determine what control actions are necessary to keep the system operating effectively. For example, the mental model of a pilot may

include information about the physical aircraft, such as the attitude and the state of the control surfaces. If the aircraft has control automation (true for almost all aircraft today), the pilot has to have information about how that automation works and the current state of the automation such as the operating mode. The automation also has a model of the process state, such as the current state of the physical control structures.

Accidents in complex systems often result from inconsistencies between the model of the process used by the controller and the actual process state, which results in the controller providing unsafe control actions. For example, the autopilot software thinks the aircraft is climbing when it really is descending and applies the wrong control law; a military pilot thinks a friendly aircraft is hostile and shoots a missile at it; the software thinks the spacecraft has landed and turns off the descent engines prematurely; the early warning system thinks the country has been targeted by someone with hostile intent and launches an interceptor at a friendly target. Note that it does not matter whether the incorrect process model was a result of an unintentional or intentional cause, which means that security can be handled in the same way. The Stuxnet worm in the Iranian reactor program is an example. The worm made the controller's process model think that the centrifuges were spinning slower than they were. The controller reacted by sending "increase speed" commands to the centrifuges, wearing them out prematurely and slowing down the Iranian development of nuclear weapons.

Part of the challenge in designing an effective safety control structure is providing the feedback and inputs necessary to keep the controller's model of the controlled process consistent with the actual state of the controlled process. An important component in understanding why accidents and losses occurred involves determining how and why the controls were ineffective in enforcing the safety constraints on system behavior. Often this is because the process model used by the controller was incorrect or inadequate in some way.

To summarize: Using systems theory, an accident results not from failures of system components, but by the inability of the safety control structure to maintain constraints on the behavior of the system. For example, there is insufficient lift to keep the aircraft airborne or insufficient distance from fixed terrain is maintained to avoid a collision. Physical failures may occur and lead to inadequate control, but there are many more things to consider.

To apply system theory to safety, a new accident causality model is required that extends what is currently used. STAMP (System-Theoretic Accident Model and Processes) expands the traditional model of causality beyond a chain of directly-related failure events or component failures to include more complex processes and unsafe interactions among system components.[13] In STAMP, accidents are caused by complex interactions among physical systems, humans, and social systems. Safety is treated as a dynamic control problem rather than a failure prevention problem. No causes are omitted from the STAMP model (event-chain models are a subset), but more are included and the emphasis changes from preventing failures to enforcing constraints on system behavior.

Some advantages of using STAMP are:

- It applies to very complex systems because it works top-down from a high level of abstraction rather than bottom up.
- It includes software, humans, organizations, safety culture, etc. as causal factors in accidents and other types of losses without having to treat them differently or separately.

STAMP is the accident causality model that underlies CAST.

[13] See Nancy G. Leveson, *Engineering a Safer World*, MIT Press (2012), Cambridge MA.

In summary, the accident causality model and the information used to analyze why a loss occurred, will have a major impact on the causal factors identified and on the conclusions reached in an accident investigation. It's now time to introduce CAST, starting with its goals.

Goals for an Improved Accident Analysis Approach

To summarize what has been presented so far**,** identifying what is labeled a "root cause" or focusing on blame unnecessarily limits what is considered in an accident investigation. Too many factors are missed and never fixed. It also limits the scope and effectiveness of the investigation as, after a loss, everyone tries to point fingers at someone else.

Instead, the goal of the causal analysis should be to learn from a loss in order to prevent future ones—not to find someone or even something to blame. Many of these people, usually low-level operators, have simply been caught up in a system where a loss was inevitable and that person was unlucky enough to be the one that triggered or was involved in the events for an accident that was "waiting to happen."

An accident analysis technique that maximizes learning from accidents and incidents should have the following goals:

1. Include all causes (optimize learning) and not focus on a few so-called "root" or "probable" causes.
2. Reduce hindsight bias.
3. Take a system's view of human behavior.
4. Provide a blame-free explanation of why the loss occurred; consider "why" and "how" rather than "who."
5. Use a comprehensive accident causality model that emphasizes why the controls that were created to prevent the particular type of loss were not effective in the case at hand and how to strengthen the safety control structure to prevent similar losses in the future.

CAST was created with these goals in mind and was developed and refined from experience gained over several decades of investigating complex accidents. The rest of this handbook will show you how to do a CAST analysis.

Chapter 4: Performing a CAST Analysis

Basic Components of CAST

CAST (Causal Analysis Based on Systems Theory) is a structured technique to analyze accident causality from a systems perspective.[14] CAST is an analysis method, not an investigation technique. But performing the CAST analysis as the investigation proceeds will assist in identifying what questions need to be answered and what information needs to be gathered during the investigation in order to create a comprehensive explanation as to why the loss occurred and to help formulate recommendations to prevent related accidents in the future.

Because the cause of an accident is defined in STAMP to be a safety control structure that did not prevent the loss, then the goal of the accident investigation is to identify why the safety control structure was unable to enforce the safety constraint that was violated and to determine what changes in the control structure are required to prevent a related loss in the future. In most cases, investigators will find that there was inadequate control provided at all levels of the safety control structure, not just the lower levels.

CAST has five parts:

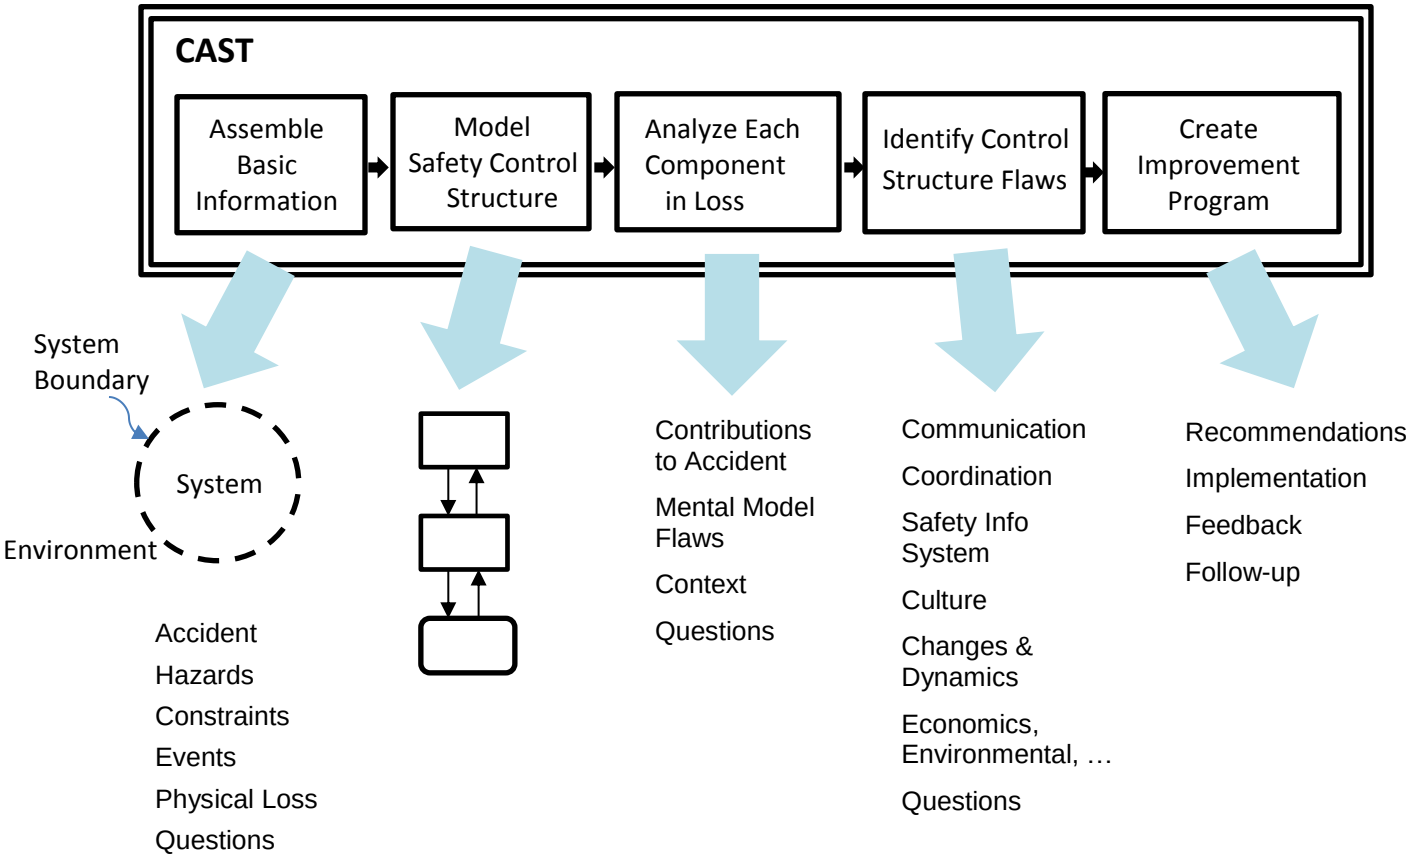

[14] It is unfortunate that the same acronym is used for Commercial Aviation Safety Team. Both uses have existed for so long that there does not seem to be a simple solution. In this handbook, CAST is always used to mean Causal Analysis based on System Theory.

Basic Components of a CAST Analysis

1. *Collect the basic information to perform the analysis:*
   a. *Define the system involved and the boundary of the analysis,*
   b. *Describe the loss and hazardous state that led to it*
   c. *From the hazard, identify the system-level safety constraints required to prevent the hazard (the system safety requirements and constraints).*
   d. *Describe what happened (the events) without conclusions nor blame. Generate questions that need to be answered to explain why the events occurred.*
   e. *Analyze the physical loss in terms of the physical equipment and controls, the requirements on the physical design to prevent the hazard involved, the physical controls (emergency and safety equipment) included in the design to prevent this type of accident, failures and unsafe interactions leading to the hazard, missing or inadequate physical controls that might have prevented the accident, and any contextual factors that influenced the events.*

   The goal of rest of analysis is to identify the limitations of the safety control structure that allowed the loss and how to strengthen it in the future.
2. *Model the existing safety control structure for this type of hazard.*
3. *Examine the components of the control structure to determine why they were not effective in preventing the loss: Starting at the bottom of the control structure, show the role each component played in the accident and the explanation for their behavior (why they did what they did and why they thought it was the right thing to do at the time).*
4. *Identify flaws in the control structure as a whole (general systemic factors) that contributed to the loss. The systemic factors span the individual system control structure components.*
5. *Create recommendations for changes to the control structure to prevent a similar loss in the future. If appropriate, design a continuous improvement program for this hazard as part of your overall risk management program.*

These are not rigid steps or a straight-line process. Work on each of the parts may proceed throughout the investigation as deemed appropriate and practical, although the first two parts will provide the basic information for the later activities and probably need to be at least started before attempting the other parts. As more is learned during the investigation, analyses will be revisited and results changed or holes filled in.

Each step involves generating questions to answer later as more is learned. The questions will help the investigators determine what more needs to be learned to explain in depth why the loss occurred. The goal at the end of the investigation is to be able to answer all the questions or to determine that they are unanswerable. The answers to the questions will provide the “why’s” for the events.

A required format for recording the results of the analysis is not provided in this handbook. Different formats may be appropriate in different situations. Some examples are provided at the end of the chapter but you may find a better format for your purposes. In addition, some industries, such as aviation, have a required format and contents for the final investigation report, which may influence the format for the CAST analysis itself.

A running example analysis is used in this handbook to explain the CAST process. It is a real accident: an explosion of a chemical reactor and a fire on June 3, 2014 at the Shell Moerdijk plant in the Netherlands. The contents of the reactor were released into the wider environment, while sections of the reactor itself were blasted across 250 meters and other debris was later found 800 meters away.

The explosion could be heard 20 kilometers away. Two people working opposite the exploding reactor were hit by the pressure wave of the explosion and the hot and burning catalyst pellets that were flying around. A large, raging fire occurred, generating considerable amounts of smoke.

Figure 12: The Shell Moerdijk Explosion

Links to other online examples of real CAST analyses on a large variety of accidents, including the running example used here, is provided in Appendix A.

The complete CAST analysis for this accident is too large to include in this handbook but can be accessed online through the link provided in Appendix A. Technical information about the chemical process that goes beyond what is needed to understand the CAST example is provided in Appendix B, along with a summary of the full CAST analysis results.

The rest of this chapter describes how each of the CAST process steps is performed on the example accident. I recommend that you take an accident with which you are familiar and go through the CAST steps yourself on that accident as you read the explanation and the examples.

Assembling the Foundational Information

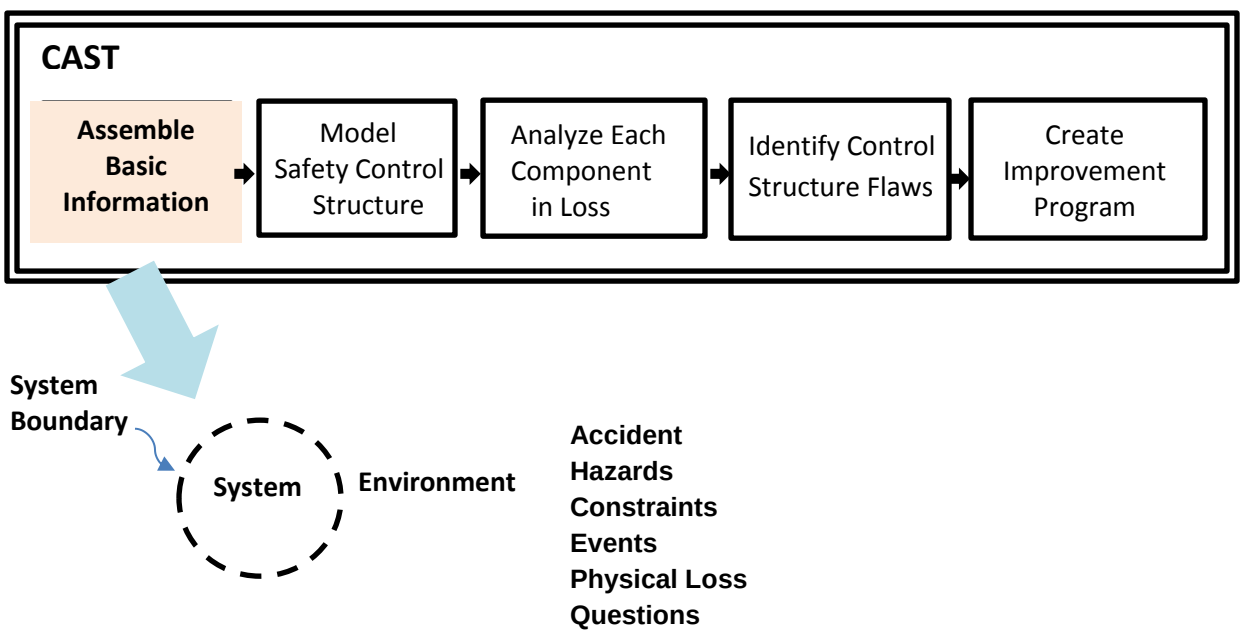

1. Collect the basic information to perform the analysis:
   a. Define the system involved and the boundary of the analysis,
   b. Describe the loss and hazardous state that led to it
   c. From the hazard, identify the system-level safety constraints required to prevent the hazard (the system safety requirements and constraints).
   d. Describe what happened (the events) without conclusions nor blame. Generate questions that need to be answered to explain why the events occurred.
   e. Analyze the physical loss in terms of the physical equipment and controls, the requirements on the physical design to prevent the hazard involved, the physical controls (emergency and safety equipment) included in the design to prevent this type of accident, failures and unsafe interactions leading to the hazard, missing or inadequate physical controls that might have prevented the accident, and any contextual factors that influenced the events.

The goal of rest of analysis is to identify the limitations of the safety control structure that allowed the loss and how to strengthen it in the future.

The first step is simply to gather basic information about what happened and to identify the goals for the analysis. The system hazard that occurred and the safety constraint violated are first identified. While seemingly simple and obvious, the hazard(s) and safety constraint(s) are important in identifying the controls included in the safety control structure in order to prevent this hazard and enforce the constraints. Starting from all constraints and controls is too inefficient.

We start by identifying the boundaries of the system of concern. In the Shell Moerdijk accident, the system analyzed is the chemical plant, its equipment, and the workers in the plant as well as the public in the area around the plant. The loss involved an explosion of Unit 4800 during startup after a maintenance stop. More details are included in Appendix B.

Next, the hazards that led to the loss and the constraints that must be satisfied in the design and operation of the system are identified.

System Hazard 1: Exposure of the public or workers to toxic chemicals.

Safety Constraints:

1. Workers and the public must not be exposed to potentially harmful chemicals.
2. Measures must be taken to reduce exposure if it occurs.
3. Means must be available, effective, and used to treat exposed individuals inside or outside the plant.

System Hazard 2: Explosion (uncontrolled release of energy) and/or fire.

Safety Constraints:

1. Chemicals must be under positive control at all times, i.e., runaway reactions must be prevented.
2. Warnings and other measures must be available to protect workers in the plant and to minimize losses to the outside community.
3. Means must be available, effective, and used to respond to explosions or fires inside or outside the plant.

In this case, there were two hazards that occurred. The constraints include the concerns of the investigation and the safety control structure to be considered that extend beyond the boundaries of the plant itself and, indeed, beyond the responsibilities of Shell. For example, responsibility for public (community) health is not normally the sole responsibility of the owners of the chemical plant itself, although they may by law be required to participate.

Deriving the safety constraints from the hazard is rather obvious, except perhaps for the inclusion of constraints to handle the case where the hazard is not avoided. In System Hazard 1 for the chemical plant (*Exposure of the public or workers to toxic chemicals*), the first safety constraint is simply a translation of the hazard to a goal statement, i.e., workers and the public must not be exposed to potentially harmful chemicals.

The second and the third are more easily missed. There is always a possibility that the safety constraint is not enforced. In that case, the hazard cannot be avoided, but it often can be mitigated and any potential harm reduced if the hazard occurs. That is the purpose of the second and third safety constraints, which require that the designers and controllers of the system and of the system's environment respond appropriately (1) to mitigate the impacts of the hazard as much as possible if the hazard does occur and (2) to minimize any potential losses if the hazard's impact cannot be totally mitigated.

Specifying the hazard itself is a bit trickier. In Chapter 1, a hazard is defined as something that has to be under the control of or within the boundaries of the system being designed. That argument is not repeated here. In engineering, designers and operators can only prevent something that they have control over. In addition, design and operations have to be performed considering the worst-case environment, not the expected or average case.

A common mistake is to state the hazard with respect to a system component and not as a statement about the system as a whole. Failure or unsafe behavior of a system component is not a system hazard—it is a cause of a hazard. Causes will be identified later in the CAST process. By starting at the system level and later generating the causes of the system hazards, you are much less likely to omit whole pieces of the puzzle in your causal analysis. It will also be easier to identify when causes are missing.

Here are some examples of proper system hazards and some that are not:

<u>System hazard</u>: The aircraft stalls and does not have adequate airspeed to provide lift.

Non-system-hazard: The engines do not provide enough propulsion to remain airborne.

Non-system- hazard: The pilots do not keep the angle-of-attack down to prevent or respond to a stall.

<u>System hazard</u>: The automobile violates minimum separation requirements from the car ahead

Non-system-hazard: The driver maneuvers the car in a way that violates minimum separation

Non-system-hazard: The automation does not slow the car adequately to avoid violating minimum separation from the car ahead

Non-system-hazard: Brake failure.

<u>System hazard</u>: Explosion and fire in a chemical plant

Non-system-hazard: Failure of the pressure release valve

Non-system-hazard: Over-pressurization of the reactor

Non-system-hazard: Operators not maintaining control over the reactor pressure

Non-system-hazard: Inappropriate reactor design for the chemical being produced.

Why are these limits on the specification of hazards being imposed? Consider the last example of an incorrect system-level hazard: "inappropriate reactor design for the chemical being produced." The problem is that it ignores the question of whether that chemical should have been produced at all or whether it should have been produced under the larger plant conditions that existed at the time. Instead, it focuses on design of the reactor and not on the operations of the other components of the plant, which may have contributed to the loss. It ignores management decisions that may have played a role. Instead, it immediately focuses the investigation down to the physical design level of part of the system, that is, the reactor design itself.

Remember, our goal is not just to find an explanation for the events but to identify the most comprehensive explanation and *all* contributions to the loss in order to maximize learning and to prevent as many future losses as possible. Focusing on the behavior of one component will lead to missing the contributions of other system components and the interactions among multiple components, like the car, the driver, the design of the interface between the car and the driver, etc. Rarely, if ever, is only one system component involved in a loss; focusing too narrowly leads to missing important information.

At this point, the events can be delineated if desired. Given the limited usefulness of the chain of events in investigating a loss, generating it is not required to complete a CAST analysis. It can, in fact, divert attention away from where it should be directed. Focusing on events alone does not provide the information necessary to determine why the events occurred, which is the goal of the CAST analysis.

While not required to start a CAST analysis, identifying the proximate events preceding the loss may sometimes be useful in starting the process of generating questions that need to be answered in the accident investigation and causal analysis. Table 1 shows the primary proximate events leading up to and following the explosion at Shell Moerdijk and some questions they raise that any accident analysis should answer. Once again, the goal is to identify ways to prevent such accidents in the future, not to find someone or something to blame.

Remember that the goal of listing the events is NOT to select one or several to identify as the cause of the loss. Instead the goal is to generate questions for the investigation that will be used in the overall causal analysis. Do not use blame-laden words such as the operators *failed* to … or "should have" in describing the events. These imply conclusions and blame. Simply describe what the hardware or operators did or did not do.

Table 1: Proximal Events Leading up to the Shell Moerdijk Loss

ID	Event	Example Questions Raised
1.	The plant had been shut down for a short, scheduled maintenance stop (called a pit stop) to replace the catalyst pellets and was being restarted at the time of the explosion	*Accidents usually occur after some type of change (planned or unplanned). The change may commonly involve a shutdown, a startup, or maintenance (including a workaround or temporary "fix"). Was there an MOC (Management of Change) policy for the plant and the company? If so, was it followed? If it was not followed, then why not? If it was followed, then why was it not effective?*
2.	One of the restart procedures is to warm up the reactors with ethylbenzene. During the warming (reheating) process, uncontrolled	*Why were the reactions unforeseen? Were they foreseeable? Were there precursors that might have been used to foresee the reactions? Did*

	energy was released and unforeseen chemical reactions occurred between the warming up liquid (ethylbenzene) and the catalyst pellets that were used.	*the operators detect these reactions before the explosion? If not, then why not? If they did, why did they not do anything about controlling them?*
3	The reactions caused gas formation and increased pressure in the reactors.	*Did the design account for the possibility of this increased pressure? If not, why not? Was this risk assessed at the design stage?*
4	An automatic protection system was triggered that was designed to prevent liquid from entering the exhaust gas system (flare). But preventing the liquids from entering the flare also prevented the gases in the system from being discharged, increasing pressure in the reactor.	*Did the operators notice this? Was it detectable? Why did they not respond? This seems like a predictable design flaw. Was the unsafe interaction between the two requirements (preventing liquid from entering the flare and the need to discharge gases to the flare) identified in the design or hazard analysis efforts? If so, why was it not handled in the design or in operational procedures? If it was not identified, why not?*
5	Continued warming up of the reactors caused more chemical reactions to occur between the ethylbenzene and the catalyst pellets, causing more gas formation and increasing pressure in the reactor.	*Why wasn't the increasing pressure detected and handled? If there were alerts, why did they not result in effective action to handle the increasing pressure? If there were automatic overpressurization control devices (e.g., relief valves), why were they not effective? If there were not automatic devices, then why not? Was it not feasible to provide them?*
6	The pressure rose so fast that it could no longer be controlled by the pressure relief devices, and the reactor exploded due to high pressure and the separation vessel collapsed and exploded.	*Was it not possible to provide more effective pressure relief? If it was possible, why was it not provided? Was this type of pressure increase anticipated? If it was anticipated, then why was it not handled in the design or operational procedures? If it was not anticipated, why not?*
7	The contents of the reactor and its associated separation vessel were released into the wider environment. Sections of the reactor were blasted across 250 meters while other debris was later found 800 meters away. The explosion could be heard 20 kilometers away.	*Was there any way to contain the contents within some controlled area (barrier), at least the catalyst pellets?*
8	Two people working opposite Unit 4800 at the time of the explosion were hit by the pressure wave of the explosion and the hot and burning catalyst pellets that were flying around.	*Why was the area around the reactor not isolated during a potentially hazardous operation? Why was there no protection against catalyst pellets flying around?*
9	A large, raging local fire occurred, generating considerable amounts of smoke.	

10 Community firefighting, healthcare, crisis management, and crisis communications were initiated.

Notice that the word "failed" does not appear anywhere in the events description (and will not appear anywhere in the CAST analysis unless a physical component failure occurred). Nor are hindsight bias statements made such as "the operators should have …." It is way too early to start making judgments and conclusions (which should be avoided even later). In addition, the questions identified in this early part of the analysis are very general. As more is learned, these questions will be refined and the answers will generate many more questions. At the end, the goal is to answer all the questions or to determine that they cannot be answered. The answers together will provide an in-depth understanding of why this accident occurred and why the controls and protection devices did not mitigate the effects of the events.

The example event chain provided here is far from "complete" (whatever that might mean because more events could always be added). In fact, their completeness will have little impact on the results of the CAST analysis. The questions that need to be answered will be generated later if not here. The event chain simply can provide a starting place if one is not obvious. Again, starting with an event chain is not strictly necessary when doing a CAST analysis.

SUGGESTED EXERCISE: Take an accident with which you are familiar or for which you have access to a detailed description of what happened. Write down the major events. The list need not be complete; this is only a starting point. Create questions associated with the events that need to be answered to understand why the events occurred. Your questions will again almost surely not be complete unless the accident is very simple. The questions you create simply form a starting point in your thinking. Many more will be created (and hopefully answered) as you go through the CAST process.

Understanding What Happened in the Controlled Process

Identifying "why" something occurred requires first knowing "what" occurred. Physical injury or losses necessarily involve a physical process while other types of losses (e.g., financial) may involve a non-physical process. In either case, the first step is to understand what happened in the controlled process. While CAST can be (and has been) applied to losses involving other types of processes, such as business processes or even criminal and legal processes, the examples in this chapter focus on physical processes. The next chapter provides examples of applying CAST to social systems.

Explosions and fires are well-known hazards in chemical plants. There are usually a large number of protection systems, alarms, sensors, pressure relief valves, and so on. We start by examining the physical controls at the plant to determine why they did not control the explosion. At this point, a CAST analysis will not differ significantly from that done in most accident analysis except that in a CAST analysis, more than just physical failures are considered. For Shell Moerdijk, at the physical level, the CAST analysis would start by generating the following description of the physical controls:

Requirements for hazard mitigation:

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. Provide feedback about the state of the safety-critical equipment and conditions
4. Provide indicators (alarms) of the existence of hazardous conditions
5. Convert released chemicals into a non-hazardous or less hazardous form
6. Contain inadvertently released toxic chemicals
7. Provide physical protection against human or environmental exposure after release

Controls:

The physical safety equipment (controls) in a chemical plant are usually designed as a series of barriers to satisfy the above hazard mitigation requirements. The Shell Moerdijk plant had the standard types of safety equipment installed. Not all of it worked as expected, however.

Emergency and safety equipment related to this accident were:

1. An automatic protection system to release gas to flare tower
2. Pressure relief devices in case of overpressurization
3. Alarms
4. Temperature sensors in the reactor

Next, we determine what happened. What physical failures and interactions led to the hazard?

Failures: None of the physical controls failed except for the final collapse of the reactor and separation vessel after pressure reached a critical level.

Unsafe Interactions: Accidents often result from interactions among the system components. In this case, the following unsafe (and mostly unexpected) interactions occurred:

1. The process to distribute the ethylbenzene over the catalyst pellets (wet them) resulted in dry zones. There were two main reasons for these dry zones:
   - The nitrogen flow was too low. To wet the catalyst properly, an adequate amount of ethylbenzene and nitrogen in the correct ratio must pass through the distribution plate. Because the flow of nitrogen was too low, the distribution plate did not operate properly. Later, due to this problem, along with other unintended interactions, the pressure increased eventually to the point where it exceeded the flow of nitrogen to the reactor. The nitrogen flow came to a standstill, resulting in a negative pressure differential.
   - The flow of ethylbenzene was unstable and at times too low. In addition to a sufficiently high nitrogen flow, a constant and sufficient flow of ethylbenzene is required in order to properly wet the pellets. The two reactors of Unit 4800 have different diameters, which means that reactor 1 requires an ethylbenzene flow of approximately 88 tons per hour while reactor 2 needs approximately 22 tons per hour. A constant flow of this volume was achieved in reactor 1. A constant flow of the correct volume was also initially achieved for reactor 2. However, once ethylbenzene began being heated, the flow became unstable. In the last hour before the explosion, this flow was virtually zero on two occasions. As a result, the ethylbenzene was not

evenly spread over the catalyst pellets, leading to the catalyst pellets not being adequately wetted and dry zones developing in reactor 2.

2. Energy released during the warming of the reactor led to unforeseen chemical reactions between the warming up liquid (ethylbenzene) and the catalyst pellets in the dry zones. As heating took place, the ethylbenzene began to react with one of the catalyst elements (barium chromate), generating heat. The ethylbenzene dissipated this heat in the areas that were sufficiently wetted. In the dry zones, however, this heat did not dissipate due to the lack of ethylbenzene. The result was that in the dry zones, the catalyst pellets heated up considerably, and there was localized development of very hot areas or "hotspots." The hotspots were not automatically detected due to the limited number of temperature sensors in the reactors.
3. Due to the rising temperature, the reaction in the hotspots kept accelerating, thereby producing even more heat. The localized temperature was now very high, which resulted in a chemical reaction between the ethylbenzene and another catalyst element (copper oxide). This reaction caused gases to be released. These follow-on reactions reinforced each other and could no longer be stopped: a runaway had developed. The rapidly rising temperature led to localized ethylbenzene evaporation.
4. Gas formation increased the pressure in the reactor. At the same time, the maximum liquid level in the second separation vessel was exceeded, causing the automatic protection system (used to release excess pressure) to shut down automatically in order to prevent liquids from entering the exhaust gas system (flare). As a result, the gases in the system could no longer be discharged. This automatic protection device to prevent liquids from entering the flare operated as designed, but had the unintended consequences of preventing the venting of the gas.
5. The buildup of gas caused the pressure to increase. Eventually, the pressure reached the point where the automatic pressure relief devices in place could not adequately release it. The pressure relief devices on the separation vessels were not designed for such rapid pressure increases, and eventually the collapse pressure of the reactors was reached. Reactor 2 collapsed and exploded, followed 20 seconds later by the explosion of the first separation vessel.
6. The contents of reactor and the separation vessel spread beyond the boundary of Unit 4800. A pressure wave and hot, burning catalyst pellets hit workers in the area causing injuries.
7. There are three remote-controlled containment valves. The explosions made these valves ineffective. The alternative was to use other limiting valves, but these valves cannot be remotely operated (they must be operated manually). Due to the intensity of the fire that had broken out and the risk of explosion, it was not possible for these to be operated immediately. An initial attempt was made around 02:30.

The above section shows why one does not need to start by listing the proximal events in isolation from the CAST analysis. The important events (failures and unsafe interactions) will be identified here as well as questions arising from these events.

After addressing <u>what</u> happened, we need to look at <u>why</u> it happened. Most of the reasons will be found in the later analyses, but physical design deficiencies and contextual factors are included here.

Missing or inadequate plant physical controls that might have prevented the accident:

1. There was an inadequate number of temperature sensors in the reactor to detect hot spots.
2. The plant was not fitted with pressure relief valves that would have prevented a runaway. Those that were installed were not designed for the rapid pressure increases that occurred.

Contextual Factors:

1. The plant was out of operation for a short, scheduled maintenance to replace the catalyst-causing granules.
2. Unexpected reactions occurred due to vulnerabilities related to the design, including the:
   - Potential for insufficient wetting (hot spots) due to design flaws.
   - Use of ethylbenzene and a design assumption that this substance is inert in the presence of the catalyst pellets.

Summary of the Role of the Physical Components in the Accident: None of the physical controls failed. The final physical collapse of the reactor and separation vessel occurred after pressure reached a critical level, which resulted from unexpected and unhandled chemical and physical interactions. Many of these unsafe interactions were a result of design flaws in the reactor or in the safety-related controls.

There are some things to be noted in this description. The description is restricted to the physical system. Operators, process control systems, etc. are not—and should not be—mentioned. The role they played will be considered later.

Also, remember that the accident model (STAMP) underlying CAST says that accidents result from lack of control, which may include control of failures. In this part of the analysis, an understanding is gained about why the physical controls were unsuccessful. Controls are used to enforce constraints on the behavior of the system components and the system as a whole and the identification of the inadequate controls will assist in refining the high-level system hazards and the safety constraints needed to prevent the hazards.

The original hazards and constraints were very general and apply to almost any chemical plant process. In fact, these could be identified for an industry as a whole and used for any accidents. Identifying the specific failures and unsafe interactions that occurred in the physical process will allow the general system hazard and safety constraints to be refined for this particular loss and therefore provide increased direction for the causal analysis.

As an example, in the Shell Moerdijk case, the general constraint to keep chemicals under positive control can now be refined to identify the specific constraints that were violated and thus the specific goals for this accident analysis:

1. Chemicals must be under positive control at all times, i.e., runaway reactions must be prevented (see the general Safety Constraint 1 under System Hazard 2 on page 37).
   a. Distribution of ethylbenzene over the catalyst pellets must not result in dry spots during a reactor restart.
   b. Protection and warning systems must be created to react to the heat and pressure generated by unintended reactions related to dry spots if they nonetheless occur.

Constraints are theoretically enforced either through eliminating the related hazardous states from the physical system design, providing controls in the physical design to prevent or minimize the hazard, or providing operational controls. Usually, all are used as engineering design can almost never be perfect. So, the goal of the analysis now is to identify why the design flaws were introduced, i.e., why the attempts to enforce the constraints through physical design were ineffective and why the operational controls also were not successful.

To achieve this analysis goal, we start by looking at the safety control structure to identify its limitations in preventing the accident. Normally, both system development and system operations need to be included. Too often, accident analysis only focuses on operations and not system development and, therefore, only operational limitations are considered. Deficiencies in the way the system was developed are not always identified or fully understood.

Modeling the Safety Control Structure[15]

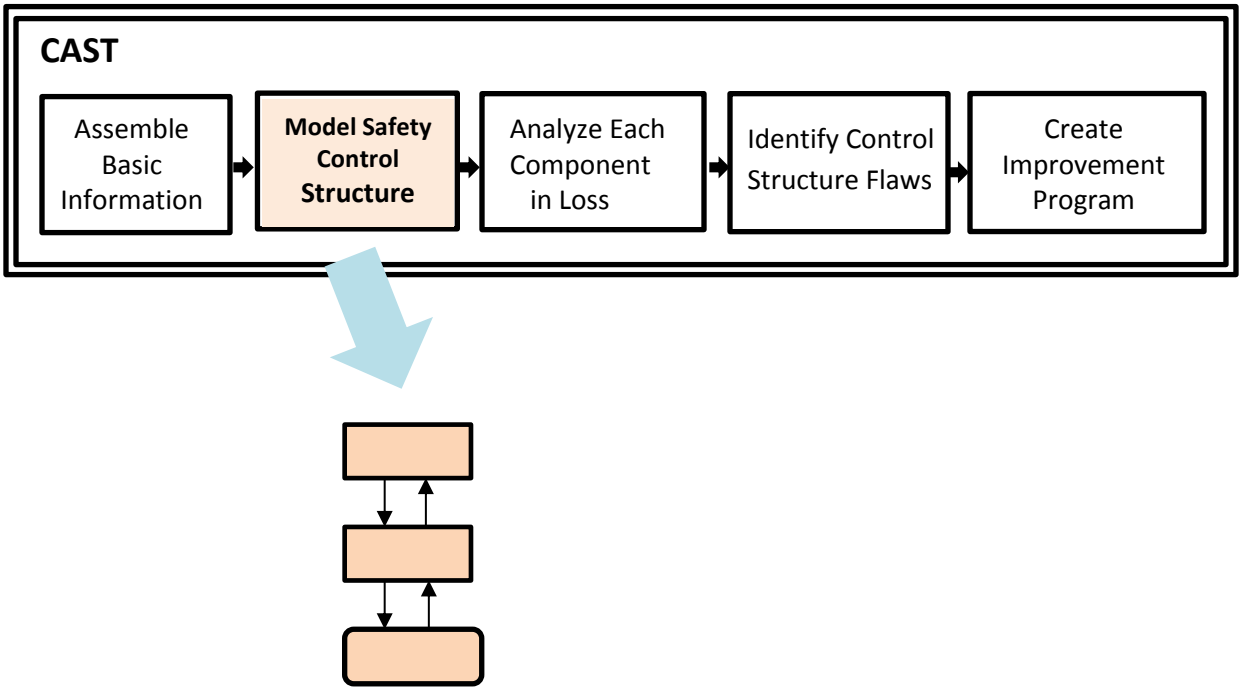

[15] A more common name for the safety control structure is the Safety Management System (SMS), although some standards for safety management systems are incomplete with respect to the safety control structure. The safety control structure includes functions and components not always in an SMS.

2. *Model the existing safety control structure for this type of hazard.*

The model of causality underlying CAST treats safety as a *control* problem, not a *failure* problem. The cause is always that the control structure and controls constructed to prevent the hazard were not effective in preventing the hazard. The goal is to determine why and how they might be improved.

Because CAST focuses on the controls and controllers and their role in the accident, modeling the control structure is necessary to start the analysis. There is no single or best way to accomplish this, assuming that a safety control structure for that system and the hazards identified in the previous step does not already exist. But there are some hints or heuristics that may help. The two most helpful are to (1) start with a very high-level, abstract control structure and then refine it later and (2) start by identifying the controls that exist for the hazard or hazards in general. These two activities may be done separately or the activities may be intertwined.

Usually, similar accidents have occurred previously, and controls have been created to prevent them. The major goal of the analysis in that case, as stated above, is to determine why they were not effective and how to improve them. Some of the controls will be common safety practices, while others may be special controls designed for unique situations. All the controls need not be identified at first. The questions generated as the analysis proceeds will identify more constraints and controls. The only mistake is to start too narrowly as this might cause you to miss important controls and controllers or system hazards. Important clues in identifying controls can usually be obtained from the original hazard analysis on the plant, assuming a hazard analysis was done that goes beyond simply identifying probabilities of component failures in order to assess the risk of the completed design.

If STPA was used for designing the system, there will be an explicit listing of the scenarios leading to an accident that were identified and the controls created during system development. If an STPA analysis for the system already exists, then it will provide a lot of information about what might have gone wrong. Theoretically, the STPA analysis should contain the scenario that occurred. If not, then there was a disconnect between the analysis during development and the operation of the system. One possibility is that the original STPA did not completely specify all the potential scenarios. Another is that the scenario that occurred was identified, but an effective control was not implemented. And, of course, the system and its environment may have changed over time after the system went into operation, negating the effectiveness of the designed controls and introducing new causal scenarios that were not analyzed originally.

If a control structure for the system does not already exist, most people find it helpful to start with a very abstract, high-level control structure that involves the general controls for this type of event. Figure 13 shows a preliminary, very high-level safety control structure for Shell Moerdijk. There are two physical processes being controlled: chemical production and public health.

The Shell Moerdijk plant has direct responsibility for control of the safety of the chemical process itself while state and local emergency management have the responsibility for preserving public health. The global Shell enterprise exerts control over the subsidiaries and Shell plants around the world. At this early point in the analysis, the type of control and responsibilities will be mostly unknown. These will need to be elaborated during the investigation. With respect to government oversight, there are local and Dutch oversight agencies and overarching EU controls. It is usually helpful to document the responsibilities (as understood so far) for each of these components with respect to the hazards involved. Again, the model will change as more information is accumulated during the investigation.

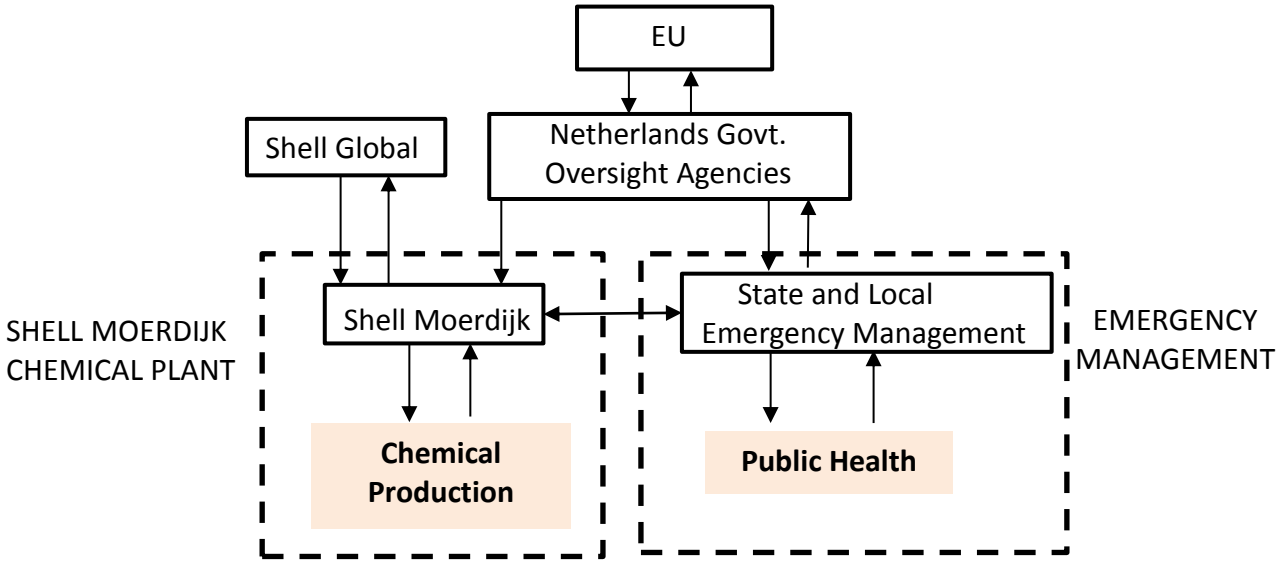

Figure 13: Very high-level safety control structure model for Shell Moerdijk.

In the process of creating the model, it may be helpful to first identify what types of controls exist both in general and in this particular system, for preventing the hazard. Once you have started the list of controls (you will probably add to it later), consider where responsibility for them exists in the physical design, human controllers, automation, or organization and social structure. For example, there will probably be physical controls (such as interlocks and relief valves), electronic controls (such as ground proximity warning systems on aircraft), electronic process control systems that monitor and adjust process parameters, human operators, equipment maintainers, management oversight, regulatory controls, and so on.

After the controls have been identified, modeling the hierarchy—to show what controls what—should be the only remaining challenge. Don't be surprised if you have some difficulty with this at first and end up making several changes. If you start with a high-level safety control structure, you should be able to place the controls in the structure as you discover them. As you proceed, more details can be added.

The control structure will probably also change as the investigation and understanding of the system develop. It may be helpful to make any changes first at the highest level that they are visible. Changes in the high-level control structure can lead to multiple changes at a lower level of abstraction. First making changes at a high-level will help to identify where extensive revisions in the control structure may be needed. For example, let's say you find late in the analysis process that maintenance was involved. Rather than try to add maintenance to the very detailed control structure developed by that time and include all the interactions it might have with other system components, it may be easier to go back to a high-level structure and add maintenance as a controller of the equipment and understand its relationship with other controllers. Once the overall role and place of maintenance in the system is modeled, more detailed control structures can be modified to include their relationships with maintenance controls.

Usually, the controls to prevent an accident are distributed around the control structure to the groups or system components best able to enforce the constraints and controls. Higher level components will have more general responsibilities, including the responsibility to oversee the operation of controls by the components below. For example, human operators may be responsible for adjusting physical parameters in the controlled process. Operations management may be responsible

for providing the training needed by the operators to perform their responsibility reliably and for hiring workers with appropriate backgrounds. Operations management probably will also be responsible for monitoring performance during operations to ensure that operators are performing adequately. Higher levels of management may ensure that training is taking place (but usually not the responsibility for specifying the details of that training) and that operations management is performing acceptably. The highest levels of corporate management and external regulatory oversight will be at yet a higher level of abstraction and detail. Together, the entire safety control structure should ensure that hazards are eliminated or controlled and safety constraints on the system behavior are enforced.

Once again, everything need not be identified at first in the causal analysis. Generating the questions and answering them as the analysis proceeds, will identify more controls and even constraints. Once several CAST analyses exist for a certain type of loss, the controls probably will have been identified in previous analyses. This step then becomes straightforward and simply involves finding unique controls for the particular loss involved. Why the controls were effective is part of the next step in the analysis. In this step, they are only listed.

Here is a demonstration of the refinement process for the Shell Moerdijk control structure, starting with the very high-level control structure in Figure 13. You need not generate a complete and final control structure at this point. Changes will likely be made and even structures added as the investigation proceeds. The goal is just to identify the controllers and their responsibilities to start the detailed CAST analysis.

Explosions and fires are well-known hazards in chemical plants, and controls are always included to prevent them. Controls will usually involve protection systems, alarms, sensors, pressure relief valves, and so on. Beyond the physical protection systems, there will usually be operators, process control systems, and safety management. Above them are various levels of management for these system components and for plant operations in general. A safety group usually exists that oversees both development and operations from a safety standpoint. This group may be responsible for various types of hazard analysis and risk assessment activities. Figure 14 includes the catalyst supplier (although it could be added later if the catalyst's role is not yet apparent) because the events point to a runaway reaction while the catalyst was being reheated. If it is not needed to explain the accident, then it can be removed later or it could be added later if omitted at this point. While the control structure can guide the analysis and make it easier to identify missing causal analysis, its major use is in the final report to help provide a coherent explanation of the accident, why it occurred, and what needs to be changed.

Responsibility for design activities rests somewhere; in the case of Shell Moerdijk, local design groups develop plant designs (not shown in Figure 14 but assumed to be under the control of and part of plant management) which are overseen or reviewed by Shell Projects and Technology at the Global level. There is also always some type of safety management at the corporate or global level. Finally, executive management (at least at the Board of Directors level) almost always has some type of corporate responsibility and oversight of safety at the enterprise management level. Here is what the control structure might look like at this point:

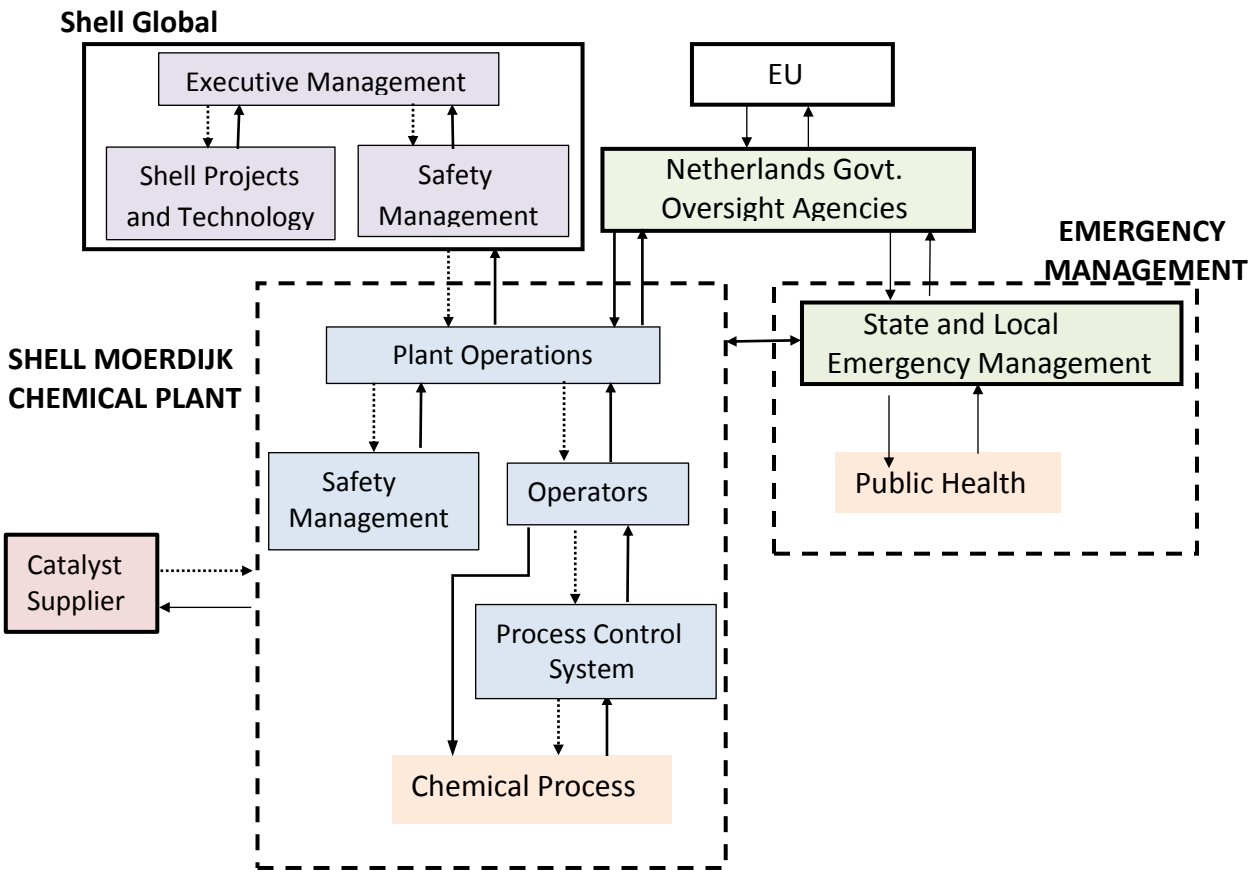

Figure 14: Shell Moerdijk safety control structure with more detail.

For readers who are familiar with STPA, which also starts with modeling the control structure, the model used in CAST need not be as detailed as that needed to perform STPA. In using CAST, most of the details will be filled in as the analysis proceeds. But the basic components should be identified at the beginning to help guide the analysis.

Once the basic components are identified, a specification of the responsibilities of each component can be started. For example, here are the responsibilities for the components included so far in the Shell Moerdijk safety control structure (Figure 14):

Process Control System safety-related responsibilities:

- Assist operators in controlling the plant during normal production and off-nominal operations (shut down, startup, maintenance, emergencies, etc.).
- Display relevant values, provide alerts, issue control actions on plant equipment.
- Control temperature, pressure, level, and flow to ensure that the process remains within the safe margins and does not end up in an alarm situation.

Operator Responsibilities:

General:

- Operate the plant in a way that does not lead to hazards
  - Monitor plant conditions and alarms
  - Control the process so that that it stays within safe boundaries of operation
  - Respond to unsafe conditions that occur.

Specific to this accident:

- Adjust gas and liquid flows as needed during startup.
- Make sure the Unit is not heated too quickly (in part to prevent damage to the catalyst pellets).

Plant Safety Management Relevant Responsibilities

- Identify plant hazards and ensure that they are eliminated, mitigated, or controlled.
- Either provide work instructions for safety-critical activities or review the work instructions provided by someone else for their safety implications.
- Ensure appropriately trained, skilled, and experienced people are assigned to high risk processes.
- Follow the Management of Change (MOC) procedures by doing a risk assessment for changes and implement risk controls based on the results.
- Provide for emergency treatment to exposed or injured individuals and ensure required medical equipment and personnel is available at all times. [*The injured personnel were treated effectively on the scene, so this aspect is not considered further.*]
- Perform audits of safety-critical activities or assist plant operations management in performing such audits [*It is not clear from the accident report who is responsible for audits, but there do appear to have been audits.*]

Operations Management relevant Safety-Related Responsibilities

- Establish safety policy for operations
- Ensure that Safety Management is fulfilling their responsibilities and providing realistic risk and hazard assessments.
- Use the results of the hazard and risk analyses provided by Safety Management in decision making about plant operations.

- Create a Shell Moerdijk Safety Management System consistent with the overall Shell Global Safety Management System and make sure it is both effective and being followed.

More specific operations management safety-related responsibilities include the following:

- Provide appropriate training for operators for nominal and off-nominal work activities.
- Follow MOC (Management of Change] procedures that require performing a risk assessment for changes or ensure that safety management is doing so. Use the risk assessment to provide oversight of the process and to design and implement risk controls in the plant and the operating procedures.
- Prepare (or at least review) the work instructions. Ensure they are safe and are being followed.
- Minimize number of personnel in the vicinity (at risk) during high-risk operations, such as a turnaround
- Keep records of incidents and lessons learned and ensure they are communicated and used by those that need to learn from them.
- Provide personnel assignments that are commensurate with the experience and training required for the activity.
- Provide a process control system that can assist operators in performing critical activities.
- Conduct audits. Establish leading indicators to be used in the audits (and in other feedback sources) or ensure that safety engineering is identifying appropriate leading indicators.

Shell Corporate Projects and Technology (Engineering) Safety-Related Responsibilities

- Create a safe design: Perform hazard analysis (or use the results of hazard analysis created by another group) and eliminate or mitigate the hazards in the design.
- Provide design, hazard, and operating information to the plant operators to assist those who are operating the plants in avoiding any hazardous scenarios that the designers were not able to eliminate or adequately mitigate in the design itself.
- Learn from the operation of their designs and improve the designs based on this feedback.

Shell Corporate Safety Management Relevant Responsibilities

- Safety of plant design, including conduct of hazard analysis on designs licensed to subsidiaries.
- Oversight of operational safety at the various Shell plants and facilities.
- Management of change procedures related to safety: creating them, making sure they are followed, and improving them using feedback from incidents.
- Communication among separate plants in different parts of the world about incidents, lessons learned, etc.
- Creating and updating a Shell-wide Safety Information System and ensuring the information is being communicated adequately both within Shell Corporate and globally and that it is complete and usable.

Executive-Level Corporate Management Responsibilities:

- Ensure that measures are taken to
  - Prevent major accidents from occurring and,
  - If accidents do occur, mitigating their consequences for humans and the environment
- Ensure the health and safety of the employees in relation to all aspects associated with the work (based on the Working Conditions Act and other regulations)
- Follow the government regulations in the countries where their plants are located.
- Create an effective safety management system and establish a strong safety culture policy. Ensure that the SMS and safety policies are being followed and they are effective.

Catalyst Manufacturer Safety-Related Responsibilities

- Provide information to customers necessary to evaluate the use of their catalyst in the reactor being designed and/or operated
- Alert customers when changes are made in the catalyst that could potentially affect the safety of its use.

Dutch Regulatory Authorities

All Dutch oversight safety and environmental authorities are grouped together here.

There are two main policies underlying this regulatory oversight:

1. Brzo: Companies must take all measures to prevent accidents and, if they occur, mitigate their consequences for humans and the environment. The company must implement this obligation by laying down policy assumptions in the Prevention Policy for Serious Accidents (PBZO), drawing up a safety report (VR), and organizing a safety management system.
2. Wabo: Regulators must check whether the company complies with regulations connected to the environmental permit, i.e., environmental safety.

General Relevant Safety-Related Responsibilities:

- Responsible for supervision and enforcement of Dutch laws to protect the environment and the public. Perform Brzo inspections focusing on process safety and Wabo inspections focusing on environmental safety.
- Responsible for enforcement of EU health and safety laws within the Netherlands.

More Specific Responsibilities:

- Identify shortcomings at companies they are responsible to oversee.
- Encourage companies to improve their safety-critical processes through supervision and enforcement. Identify shortcomings and persistently question companies to prompt them to investigate and detect deep-seated causes of incidents and accidents. Ensure that any shortcomings identified are corrected.
- Assess modifications made to plants, procedures, and processes (although they are not expected to perform the risk analyses for the companies).
- Pay greatest attention to safety-critical processes, including maintenance and reactor start-up

Emergency Services Responsibilities

- Firefighting, crisis management, crisis communications including among other things:
  - Informing citizens of the incident,
  - Measuring substances released on a coordinated basis,
  - Operating a telephone advisory line,
  - Informing citizens about the results of the measurement of the substances released and the ensuing recommendations.

Again, the control structure model is not fixed but can and most likely will evolve and change as more is learned about the accident and about the responsibilities of each of the control structure components. The CAST analysis will essentially involve identifying whether the responsibilities were effectively implemented and, if not, why not and what changes might be made to prevent this type of accident in the future. Another potential conclusion might be that responsibilities are not adequately identified or

appropriately assigned and that the basic safety control structure needs to be redesigned. Most accident investigations will likely find both types of contributors to the loss.

How does the analyst identify the responsibilities? Most of these will be documented somewhere, either in company manuals or government documents. In addition, questioning the controllers about their own and other's responsibilities during the investigation will at least elicit what the controllers think they are. Confusion and inconsistencies here could be a clue as to why the accident occurred. One of the important characteristics of the U.S. Navy's SUBSAFE program, which has been spectacularly successful over its 50+ year existence, is that everyone in the program knows exactly what are their own responsibilities and also what everyone else's are. By continually reinforcing and auditing this knowledge about responsibilities, accidents are reduced.

Individual Component Analysis: Why were the Controls Ineffective?

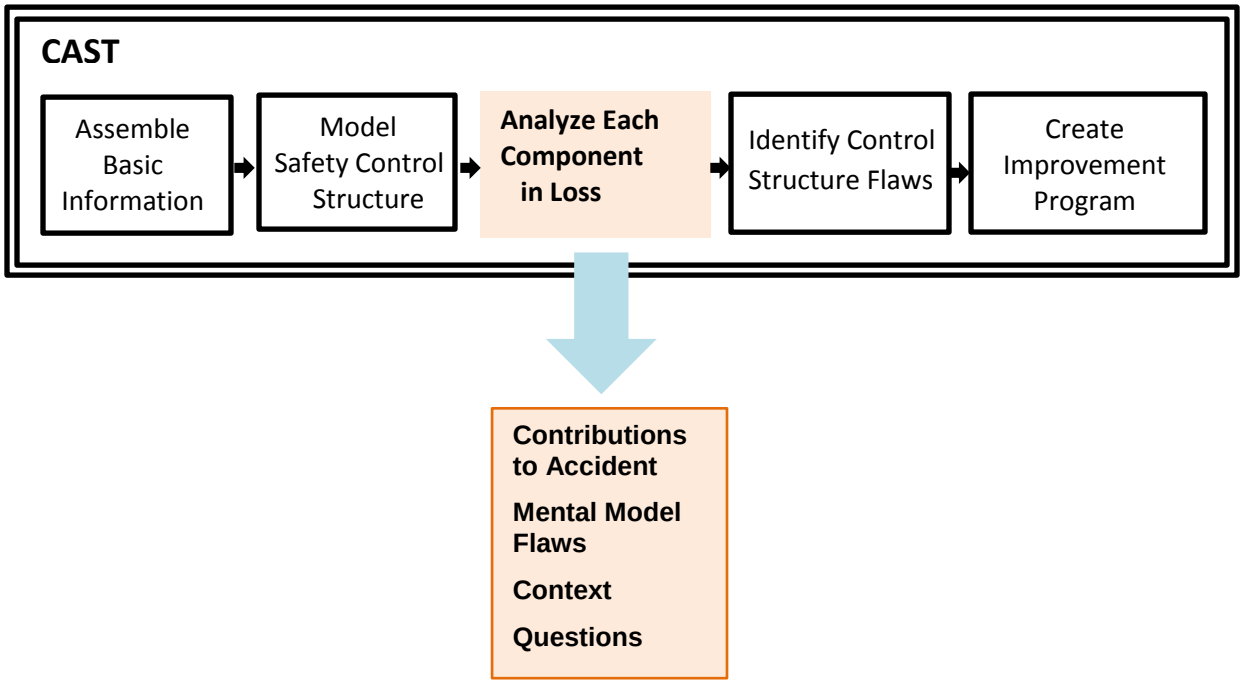

3. *Examine the components of the control structure to determine why they were not effective in preventing the loss:*

   Starting at the bottom of the control structure, show the role each controller played in the accident and the explanation for its behavior: why each control component did what it did and why the controller thought it was the right thing to do at the time.

Once the basic control structure and controls are identified, the next step is to show why the control structure, i.e., the current controls, did not prevent the accident. There will be two parts to this process. The first, described in this section, looks at the individual controllers (which may be automated or human) and the role they played in the accident. The second, described in the next section, looks at the operation of the control structure as a whole and the interactions among the components. Remember, the goal is to identify flaws in the control structure: not to assign blame to individuals or individual components but to understand why it made sense for them to act the way they did.

Also, as started in the previous steps, the CAST process includes generating questions that need to be answered during the investigation so that at the end, a complete explanation for why the accident occurred can be provided. I find it most helpful to start at the bottom components in the control structure and work my way upward.

There are several parts in the CAST analysis of each controller:

- Component responsibilities related to the accident
- Contribution (actions, lack of actions, decisions) to the hazardous state: [16]

Why?

- Flaws in the mental/process model contributing to the actions:
- Contextual factors explaining the actions, decisions, and process model flaws:

The first two parts simply document the role the component played in the loss, if any. The component safety-related responsibilities were identified when modeling the safety control structure. Only those related to the loss need to be considered. The role or contribution (behavior related to the hazardous state) will be the responsibilities that were not fulfilled. For example, one responsibility of the operators during reheating of the reactor at Shell Moerdijk after a maintenance stop is to control the process so that it stays within safe boundaries of operation. In this case, they did not adjust the gas and liquid flows in a way that prevented the overheating of the catalyst pellets. Remember, no judgmental or hindsight words should be used like they "failed" to adjust the flows or they "should have" adjusted the flows. Simply describe what happened in a non-accusatory way. The behavior that contributed to the loss is described in a straightforward manner that simply says what the person or group did or did not do.

The "why" part of the analysis involves explaining why the component behaved in the way it did, i.e., the contextual and other factors that made the behavior seem correct to them at the time. For example, contextual factors might include incorrect information that the controller had at the time, inadequate training, pressures and incentives to act the way they did, misunderstanding about their role or about how the controlled components worked, etc. This "why" part of the analysis will also involve answering the questions previously raised about the behavior of this component. Answering the "why" questions for a component usually raises more questions, which may need to be answered in the analysis of the higher levels of the control structure. The goal is that at the end of the CAST analysis, all the questions raised are answered. In reality, there may be remaining questions that just cannot be answered definitively. These should, of course, be documented in the final conclusions. If the people involved understand that the goal of the accident investigation and causal analysis are not to find "guilty parties,"

[16] Although in STPA the term "unsafe control actions" is used without much complaint, the use there is about hypothetical actions and not things that actually have occurred. Using that term in CAST lends a tint of blame and judgment and should be avoided. We have used various terms such as "Control actions contributing to the hazard" or contributory control actions or simply "role or contribution of the controller to the hazardous state" to avoid common pejorative terminology.

they are more likely to provide candid answers. Enlist their help as part of an explanatory process to improve safety in the future. This kind of trust, of course, needs to be established over time.

Focus on understanding "why" the controllers acted the way they did. Most controllers are trying to do the right thing. If they had evil intent (were suicidal or homicidal), of course, that needs to be identified but such behavior is rare in accidents. Usually the controllers are trying to do a good job, and there is some reason that they didn't. For automated controllers, understanding will involve looking at the design and the designers of the automation and the assumptions they made about what type of behavior was required. Almost all software-related accidents involve incorrect assumptions about the behavior required to prevent the hazards, i.e., a flawed understanding of the requirements for safe operation.

Understanding the reasons for the behavior that occurred will provide the best guidance possible in creating effective recommendations. If the accident report identifies only what the controller did wrong, we learn almost nothing required to prevent the same behavior in the future beyond telling people "don't do that." As the controllers usually had a good reason for doing it at the time, this type of instruction is not helpful. Rules and procedures are never mindlessly obeyed under all circumstances. If the designers of the system intend for operators to thoughtlessly follow written directions, then those control responsibilities should be automated. Checklists are often provided as guidance, but we usually want people to make judgments about the appropriateness of the checklist at the time of the emergency. If they did not follow the checklist when they should have, then we need to understand why.

Humans are left in systems today to make cognitively complex decisions about what needs to be done. If they make the wrong decisions, the reasons why (e.g., they had incorrect information, they were misled by the process control system output, they were not given adequate support to make safe decisions, there was no or misleading feedback about the state of the component they were controlling) will provide the information needed to redesign the safety control structure and safety controls to assist controllers in making better decisions in the future.

The CAST process is designed to avoid hindsight bias as much as possible. After the accident or incident, the behavior will appear to be clearly wrong. That's obvious and doesn't need to be stated. What is useful is to understand why—*at that time*—it was not obviously wrong to the controller. Why not? What types of information did they need to make a better decision? Did they have it? Did they understand what that information meant at the time? Were there other competing pressures that made it difficult for them to absorb the information or to process it? The more fully the context can be identified, the more effective any resulting recommendations will be in reducing future losses. The context may be physical, psychological, informational, managerial, and situational, such as workload or productivity pressures, incentives/disincentives, etc.

The notation used to record the results of this analysis step is irrelevant as long as the information is recorded and easy to understand. Some suggestions for notations we have been used successfully are included at the end of this chapter.

In this handbook, only a few examples from the CAST analyses of the Shell Moerdijk are shown. A summary of the role played by all the components is provided in Appendix B. Links to the complete CAST analysis for Shell Moerdijk as well as many other real accidents are provided in Appendix A.

An automated process control system is a standard controller for plants. Therefore, it's a convenient place to start in understanding the Shell Moerdijk explosion. The safety-related responsibilities include:

Safety-Related Responsibilities of the Process Control System:

1. Assist operators in controlling the plant during normal production and off-nominal operations (shut down, startup, maintenance, emergencies, etc.).

2. Display relevant values, provide alerts, issue control actions on plant equipment.
3. Control temperature, pressure, level, and flow to ensure that the process remains within the safe margins and does not end up in an alarm situation.

To identify its role in the loss, we can start with determining whether these responsibilities were satisfied. In this case, the process control system contributed to the accident by:

Contribution to the hazardous state:

1. The process control system did not provide the assistance required by the operators to safely control the start-up process, including automatically controlling the heating rate and other important variables.
2. The process control system did not step in to stop the process when pressure and temperature increased precipitously.
3. The process control system did not issue an automatic reset after two high-high level alarms, so the gas discharge system remained closed.

As an example of the detailed "why" analysis, i.e., the contextual factors for number 2 (i.e., the process control system did not stop the process when pressure and temperature increased precipitously), consider:

Why? (Contextual Factors Affecting the Unsafe Control)	Questions Raised
There was no emergency stop button for Unit 4800. Without an emergency stop, there was no way to safely stop the operation of the unit quickly with a single press of a button. No reason was provided for this design "flaw" in the accident report but could have been answered at the time of the investigation. Emergency stop buttons are standard in safety-critical systems.	*Was this complacency, cost, or was there an engineering reason?*
The instrument-based safety devices were designed to respond to and prevent particular conditions, but not the ones that occurred. After this accident, a safety device was added to protect Unit 4800 from an excessively high temperature due to an unwanted chemical reaction with hydrogen.	*Why were these conditions omitted? Was a hazard analysis performed that could have identified them?*
There is a containment system, which is described as "one or more appliances of which any components remain permanently in open connection with each other and which is/are intended to contain one or more substances." The valves of the containment system, however, cannot be remotely operated and must be operated manually. Due to the intensity of the fire that night and the risk of additional explosions, it was not possible for these valves to be operated immediately and, in fact, were not operated until several hours later when the fire had been extinguished.	*The containment system design was not useful in this situation. Was this known beforehand? Was it impossible to design one that can be operated remotely?*

An important part of the understanding of why a controller behaved the way it did may arise from an incorrect process/mental model. The CAST analysis involves both identifying the flaws in this model and determining why they occurred.

Process Control System Process Model flaws: The Shell Moerdijk process control system, for the most part, had correct information about the state of the process. There was some missing information about temperature, however, resulting from an inadequate number of temperature sensors provided in the reactors.

When the component analysis has been completed and documented (details not included here), a summary can be made of the role of the component in the accident along with recommendations. Any open questions that are important in the causal analysis should also be documented.

Summary of the Role that the Process Control System Played in the Accident: The process control system was not designed to provide the necessary help to the operators during a start-up or to allow them to easily stop the process in an emergency. The reason for these design decisions rests primarily in incorrect assumptions by the designers about the impossibility of the scenario that occurred. Even after previous incidents at similar plants in which these assumptions were violated, the assumptions were not questioned and revisited (see the CAST analysis of the safety information system).

More generally, the process control system was configured only for the normal production phase. There were no special automated control circuits for the heating phase after a catalyst has been replaced, which is the phase in which the problems arose. In fact, this phase is commonly known to be the most accident prone, so the decision not to provide automated help to the human operators during this phase is quite strange, but no explanation is provided in the accident report.

Unanswered questions: Why was a decision made not to provide assistance to the operators during the restart phase of operations? Why was an emergency stop capability not designed into the system?

Recommendations: The operators' knowledge and skill are most challenged during off-nominal phases, and most accidents occur during such phases and after changes are made or occur. The process control system should be redesigned to assist operators in all safety-critical, off-nominal operations (not just this restart scenario). For manual operations, the goal should be for the process control system to provide all necessary assistance to the operators in decision making and taking action and to reduce attention and time pressures (see the section on the operators).

The rest of this section shows examples from the Shell Moerdijk CAST analysis for different types of controllers. The process control system analysis provides an example of an automated controller. The other controllers for the Shell Moerdijk Chemical Plant were people or groups of people. Details, again, are omitted here for space reasons.

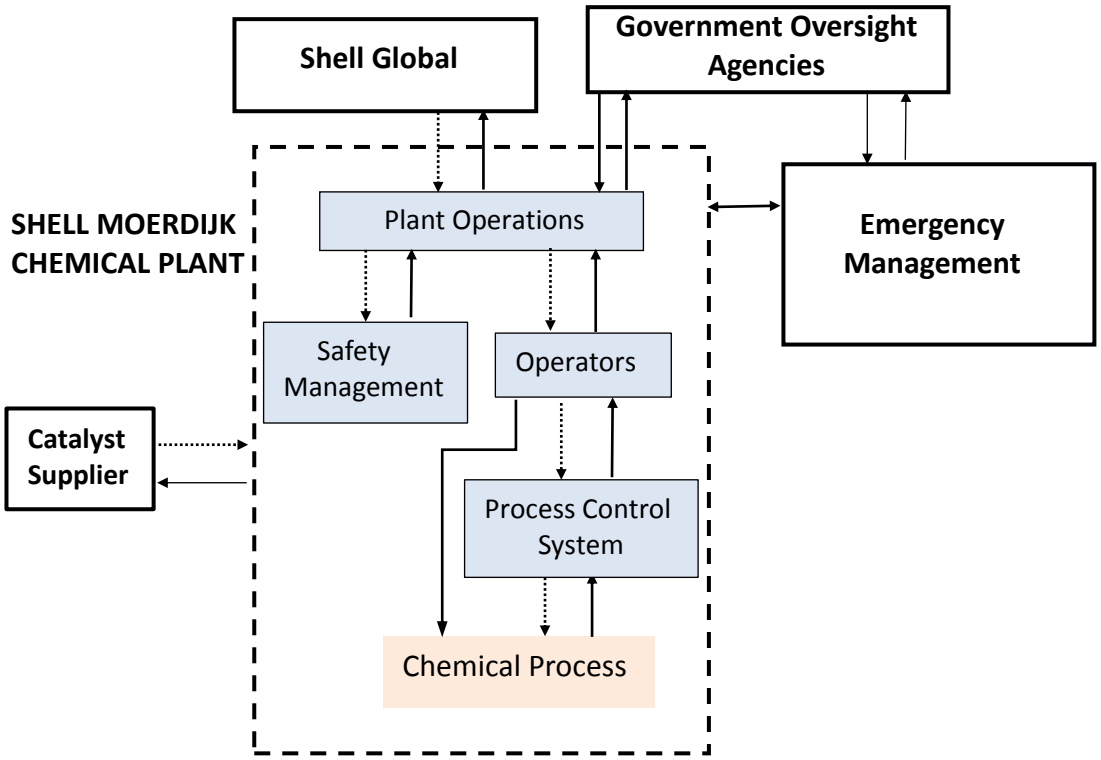

Figure 15: Shell Moerdijk Chemical Plant Safety Control Structure.

Operators

The operators' actions clearly contributed to the explosions. For example, they manually added additional warmth to the ethylbenzene at a time when heat was increasing precipitously; they did not notice and respond to hot spots and negative pressure differential; they did not respond appropriately to alarms; they left the gas discharge system closed when the gas was increasing; they did not stabilize, slow down, and stop the process when pre-set limits were exceeded (which is a fundamental principle in operator training at Shell); etc.

Given all these things that in hindsight the operators did wrong, they appear to have major responsibility for the loss. In fact, listing these actions is where many (most?) accident causal analyses and accident reports stop and the operators are deemed to be the major cause of the loss. While it is possible that everyone working the turnaround that day was negligent and irresponsible, it is more likely that they were trying to do their best. Without understanding *why* they made bad decisions, i.e. why the decisions seemed correct to them at the time, we cannot do much about preventing similar flawed decision making in the future. Many of the answers lie in higher levels of the control structure, but some of the operators' actions can be understood by looking at their process models and the context in which they were making decisions.

The relevant general safety-responsibilities assigned to the operators:

- Operate the plan in a way that does not lead to hazards
  - Monitor plant conditions and alarms
  - Control the process such that it stays within safe boundaries of operation
  - Respond to unsafe conditions that occur

And specific to this incident: • Adjust gas and liquid flows as needed during startup • Make sure the Unit is not heated too quickly (in part to prevent damage to the catalyst pellets)

None of these responsibilities were fulfilled. Examples of operator behavior that contributed to the hazardous state include:

- The operators did not stabilize or halt process before the explosion when critical process boundaries were exceeded.
- Heating was started (around 21:00) while the situation was still unstable and after only 45 minutes of wetting. Proper wetting had probably not been achieved by that time.
- The operators manually added additional warmth to the ethylbenzene at a time when heat was increasing precipitously.
- The operators did not notice and respond to hot spots. They also did not notice and respond to the related negative pressure differential.
- The operators did not properly adjust nitrogen flow. The lower (than required) nitrogen flow was one of the causes of the accident.
- The operators did not respond to alarms.

Looking only at these behaviors, the behavior of the operators seems egregious. This is where most accident reports stop. But the picture looks very different when the contextual factors and process model flaws are considered. Because of space reasons, only a few examples of these are provided below. The reader is referred to the full analysis in the CAST report for the more detailed analysis.

Why? (Contextual Factors Affecting the Unsafe Control)	Questions Raised
The operators were not provided with the information they needed to make correct decisions. The process control system provided almost no support (feedback). Throughout the accident report, there are many actions by the operator that are noted to require intense attention and deep understanding of the process dynamics.	*Why? Was any human factors analysis done when the decision was made to have the operators manually control startup to ensure that they were capable of doing this reliably? And without process control system support?*
Adjusting gas and liquid flows was known to be difficult, partly because of weaknesses in the design of the central pump. In addition, the job analysis does not provide clear instructions about how the filling and circulation has to be done. The only clear instruction was that the central pump not be allowed to run “dry” or it would break.	
Safety during the heating and wetting of the reactors was dependent on the knowledge and skill of the operators and production team leader on duty at the time. Shell Moerdijk’s written policies required that the starting and stopping of the reactors had to be done by experienced operators using the work instructions provided for this purpose. However, while the personnel performing this maintenance stop were experienced staff and were trained for working on this unit	

during regular production, a catalyst change only occurs once every three or four years and this was their first experience with restarting the reactor. At the same time, they were not provided with support from the process control system for this startup.	
The accident report says that the start-up procedures had been followed correctly by the operators. The problem, then, must have been in the procedures themselves. In fact, the work instructions used by the operators were incorrect and incomplete. The work instructions did not follow the format provided by Shell for such instructions and omitted much of the required information such as critical conditions and required steps to be taken by the operators. Nitrogen flow was not considered critical, for example, and was not included in the work instructions although it turned out to be critical after the fact.	*Why were incorrect and incomplete work instructions provided to the operators?*
The startup procedures were produced by the operators. There does not seem to have been any expert review by engineers and safety personnel. The accident report provides no explanation for *w*hy operators are writing their own work instructions for safety-critical operations without at least expert review by engineers and safety personnel.	*The report does not answer the obvious questions here: Were the work instructions produced by the operators performing the startup (who had no experience)? If it was other operators, did they have the experience and knowledge to create the work instructions? Who reviews these work instructions? Were they reviewed by anyone? (The answer to this question appears to be no).*
There is a design book, created by Shell Global Projects and Technology, that contains detailed information about the design and operation of the reactor. The Design Book was not used in creating the work instructions for the startup because, the accident report says, the book was too detailed and "intricately presented" and was not understandable by the operators charged with drawing up the work instructions.	*There are many obvious questions raised here.*

Why? (Process Model Flaws)	Questions Raised
There was a lot of hindsight bias in the accident report describing what the operators "should have" done, but the conditions that occurred were all consistent with their expectations (process model) given the conditions during previous maintenance stops. The operators considered fluctuations in pressure to be normal during restarts (process model flaw). The pressure had fluctuated continually since the beginning of the restart. Such	*Why were the expectations (their mental models of process behavior) not correct in this instance?*

fluctuations had occurred during previous restarts and that was what they expected. In hindsight, of course, these expectations turned out to be wrong. Given the situation, the number of alarms that sounded and their frequency was not out of the ordinary, without knowing *after the fact* that the process was indeed not in an ordinary state.	
The accident report says that the ability to make an assessment to intervene in special situations requires knowledge of, experience with, and thorough preparation for such special situations. The operators must fully grasp what caused the unit to exceed limits in order to assess risk. The operator did not have the required knowledge or expertise. . In fact, even the corporate safety engineers did not know that the conditions that occurred here could result in a hazardous state. If the designers did not believe such scenarios were possible, why would the operators? In fact, nobody at Shell thought that a scenario involving a fast and high-pressure buildup was possible.	*Why did the designers and corporate safety engineers believe that such scenarios were impossible?*

In summary, the actions of the operators look much less clearly negligent when the reasons why the operators behaved the way they did is the focus of the investigation rather than a focus on what the operators did that turned out to be wrong. The operators, in fact, acted appropriately or at least understandably given the context, the incorrect work instructions (which they followed), and their lack of training and required skill and knowledge in performing the work. In addition, they were provided with almost no assistance from the process control system. As many of the tasks they needed to do required intense attention, precision, mental effort, deep understanding of process dynamics, and frequent adjustments to a continually fluctuating process, such assistance would have been invaluable.

The designers of the plant did not recognize or understand the risks (see the appropriate sections of the CAST analysis) so the risks might not have been communicated thoroughly to the operators. Management seemed to rely on operators seeing something strange and stopping the process, but did not provide the information and training to ensure it was possible for operators to do this. Such a policy provides a convenient excuse to blame the operators after an accident, but it does not help the operators to carry out their responsibilities.

Generating recommendations will be covered later, but some obvious ones from this part of the analysis are that the operators must have the appropriate skills and expertise to perform their assigned activities, and there must be someone overseeing operations who is assigned the responsibility for implementing this requirement. Other possible recommendations include the need for a human factors study during the job analysis to ensure that the operators are provided with information and a work situation that allows them to make appropriate decisions under stressful conditions, better automated assistance during all phases of operation, training for activities that are known to be hazardous such as startup, and finally improved work instructions together with a better process for producing them.

This analysis, as usual, raises a lot of additional questions that need to be answered in order to understand what happened and make recommendations to prevent such occurrences in the future. These will be summarized in the rest of the section of the handbook (with the details omitted) so that readers have a fuller picture of the CAST process and the results that can be obtained from it.

So far, we have looked at the physical process, the automated process control system, and the plant operators. We now briefly look at the role that the CAST analysis identifies that each of the other controllers played in the loss, working our way up the safety control structure. Once again, detailed CAST analyses for each controller is not included in this handbook but can be found elsewhere.

Plant Safety Management

The plant safety department usually provides oversight of operational safety and provides information to plant operations management to ensure that operational decisions are made with safety in mind. Examples of the detailed analyses are included here, but we have found that too much detail discourages reading of the report. A useful compromise is to include the detail but also include a summary of the role in the component in the accident that summarizes the important points of the detailed analysis.

Relevant Responsibilities of Plant Safety Management

- Identify plant hazards and ensure that they are eliminated, mitigated, or controlled.
- Either provide work instructions for safety-critical activities or review the work instructions provided by someone else for their safety implications.
- Ensure appropriately trained, skilled, and experienced people are assigned to high risk processes.
- Follow the Management of Change (MOC) procedures by doing a risk assessment for changes and implement risk controls based on the results.
- Provide for emergency treatment to exposed or injured individuals and ensure required medical equipment and personnel is available at all times. [*The injured personnel were treated effectively on the scene so this aspect is not considered further.*]
- Perform audits of safety-critical activities or assist plant operations management in performing such audits [*It is not clear from the accident report who is responsible for audits but there do appear to have been audits.*]

Process Model Flaws of the plant safety managers

- Regarded ethylbenzene as a safe substance in this process.
- Considered the start-up process not to be high risk.
- Thought that the Operators and the Production Team Leader could manage and control the start-up manually based on their knowledge and experience.
- There were so sure that they understood the risks that even incidents at other similar plants did not trigger any doubts about their assumptions. Alternatively, they may not have been made aware of the previous incidents.

A Few Example Contextual/Process Model Factors (more included in summary below):

- Over time, understanding of the most appropriate procedures relating to Unit 4800 changed. Some of the procedures were not considered critical to safety. So, these procedures were not included (or were no longer included) in the amended work instructions.
- Safety assessment procedures complied with government requirements.
- Hazard assessment focused on processes that were considered higher risk and primarily assessed the effects of substances on the environment and not on safety.

Summary of the Role that Shell Moerdijk Safety Management Played in the Accident:

- The safety analysis methods used were either not appropriate, not applied or were applied incorrectly. However, the methods used complied with the Shell requirements and satisfied the Dutch legal and regulatory requirements. They were also standard in the petrochemical industry. Safety management did not consider certain relevant information nor investigate how ethylbenzene reacting with the catalyst could cause an explosion. Safety management at Shell Moerdijk, as is common in many places, seems to have been largely ineffectual, with lots of activity, but much of it directed to minimal compliance with government regulation. A partial explanation for their behavior is that everyone believed that a reaction between ethylbenzene and the catalyst was impossible and that the start-up process was low risk.
- Although Shell's safety management system includes requirements for dealing with changes, the Management of Change (MOC) procedures were not followed nor implemented effectively. Risks resulting from changes made to the plant, the catalyst, the processes, and the procedures were not identified and managed. Not following MOC procedures has been implicated in a very large number of accidents in all industries.
- A new catalyst was selected for use at Shell Moerdijk, but an assumption was made that the properties of the new catalyst were the same as those of the previous catalyst. The altered composition of the new catalyst was stated in the safety information sheet provided by the manufacturer, but safety engineering did not notice this change.
- Procedure changes (heating rate and nitrogen flow) were instituted without a risk assessment. These procedures were not included in the amended operator work instructions. Other plant and production changes were not systematically examined for their safety effects on the basis of a risk analysis in all cases. Only a few failure scenarios were examined, but this fulfilled both Shell Global and Dutch law. There were never any safety studies that focused on the circulation and heating of the unit involved in the accident. Safety was based on studies done in 1977 that had shown the catalyst (as it then existed) was inert in the presence of ethylbenzene. This assumption was never reassessed even though the composition of the catalyst changed over time and similar incidents occurred at other Shell installations. There were no requirements by either Shell Global or the Netherlands law to do so.
- The number of leaks was used as the primary leading indicator of process safety, which clearly has nothing to do with this accident. Because the number of leaks had been greatly reduced in recent years, decision makers assumed safety was improving. The use of this leading indicator is common in the petrochemical industry and was accepted by the Dutch regulatory authorities.
- Similar incidents occurred at Shell Nanhai and at Shell Moerdijk after initial startup in 1999. No explosion occurred at Nanhai because of special factors. These events did not trigger a response in terms of reassessing risk or procedures or the assumptions that a runaway was impossible. Recommendations following the Nanhai incident were not used to reduce risk. *Why? There is no information in the accident report about why the Nanhai incident recommendations were not implemented.*
- Safety engineering did not provide proper oversight of the generation of work instruction, which allowed unsafe work instructions to be provided to the operators. The work instructions did not follow the format provided by Shell for such instructions. They also omitted much of the required information such as critical conditions and required steps to be taken by the operators.

Recommendations:

While the problems specific to the explosions on 3 June 2014 should be fixed, there were a lot of weaknesses in the Shell Moerdijk safety management practices identified in the official Dutch Safety Agency accident report and in the CAST analysis. These need to be improved.

- Safety management at Shell Moerdijk needs to be made more effective. Safety engineering needs to be more than just going through the motions and minimally complying with standards.
- All work instructions should be reviewed for safety by knowledgeable people using information from the hazard analysis. [In this case, the hazard analysis was flawed too, but that is a different problem to fix.]
- MOC procedures must be enforced and followed. When changes occur, assumptions of the past need to be re-evaluated.
- Hazard analysis and risk assessment methods need to be improved.
- More inclusive and relevant leading indicators of risk need to be established.
- Procedures for incorporating and using lessons learned need to be established or improved.

Operations Management

Relevant Safety-Related Responsibilities

- Establish safety policy for operations.
- Ensure that Safety Management is fulfilling their responsibilities and providing realistic risk and hazard assessments.
- Use the results of the hazard and risk analyses provided by Safety Management in decision making about plant operations.
- Create a Shell Moerdijk Safety Management System consistent with the overall Shell Global Safety Management System, ensuring it is both effective and being followed.

More specific safety-related responsibilities include the following:

- Provide appropriate training for operators for nominal and off-nominal work activities.
- Follow MOC (Management of Change] procedures that require performing a risk assessment for changes or ensure that safety management is doing so. Use the risk assessment to provide oversight of the process and to design and implement risk controls in the plant and the operating procedures.
- Prepare (or at least review) the work instructions. Ensure they are safe and are being followed.
- Minimize the number of personnel in the vicinity (at risk) during high-risk operations, such as during a turnaround.
- Keep records of incidents and lessons learned and ensure they are communicated and used by those that need to learn from them.
- Provide personnel assignments that are commensurate with the experience and training required for the activity.
- Provide a process control system that can assist operators in performing critical activities.
- Conduct audits. Establish leading indicators to be used in the audits (and in other feedback sources) or ensure that safety engineering is identifying appropriate leading indicators.

Process Model Flaws of operations management personnel

- Regarded ethylbenzene as a safe substance (an inert medium) under all process conditions. Therefore, they did not consider the heating phase to be risky.
- Thought the personnel involved in the turnaround had appropriate training and experience.
- Did they not know about similar incidents at other Shell installations with the same design or did they not think they were relevant to Unit 4800?
- In general, they had an inaccurate view of the risk that existed in the plant.
- Inadequate feedback leading to these flaws:
  - The heating phase did not appear in the report provided by plant safety management.
  - Plant safety management did not provide correct risk assessments to operations management.

Contextual Factors: (omitted, see complete analysis; also included in summary below)

Summary of the Role of Operations Management in the Accident:

- Operations management did not identify the flaws in the risk analyses performed or the procedures used for these risk analyses. The risk analyses complied with the minimal requirements of the Dutch regulatory authorities and apparently with the Shell requirements.
- Changes over time were not subjected to assessment in accordance with the MOC procedures.
- Work instructions were created by the operators without safety engineering oversight. They did not comply with the required Shell format for such work instructions and did not include important criteria for the job such as heating rate. Nitrogen flow, an important factor in the accident, was ignored in the work instructions.
- A decision was made to configure the process control system to control the plant during the normal production phase but not during non-production and maintenance phases. They did not think these activities were high risk and thought that manual operation would suffice. The reasons for this decision are not in the accident report.
- Two employees from different contractors were allowed to work in the adjacent unit during the start-up, probably because they did not believe that phase was dangerous.
- Operations management did not assign operators to the start-up that had the qualifications required in the Safety Report. No reason is given in the accident report as to why this happened.
- Operations management did not ensure that lessons learned from similar plants and at Shell Moerdijk in 1999 were incorporated in the design and operation of Unit 4800.
- The Safety Management System at Shell Moerdijk did not prevent unsafe situations from being overlooked or internal procedures from not being followed. There is no information in the accident report about who created the SMS or who was responsible for ensuring that it was working properly.
- Internal Shell Moerdijk audits did not show any of these shortcomings. Not enough information is provided to determine why they were ineffective.
- Shell Moerdijk has a Business Management System in which safety management is integrated. No details are provided, but in general this is a poor practice and has been a factor in major petrochemical company accidents such as Deepwater Horizon. Decision making needs to occur with full information about all the factors that must be considered and not lost by integrating risk information in a nontransparent way.

Recommendations:

- Establish and ensure proper MOC procedures are followed. If changes occur, retest assumptions that could be affected by those changes. This implies that these assumptions must be recorded, leading indicators established for identifying when they may no longer be correct, and a process established for testing and responding to changes that might affect these assumptions.
- A thorough review of the Shell Moerdijk SMS should be done with emphasis on why it was unable to prevent this accident. Major factors in this accident are related to basic activities that should have been controlled by the SMS.
- Update procedures to eliminate the causes of the accident such as lack of control and supervision of the work instruction creation and oversight processes, inadequate hazard analysis and risk assessment procedures, assignment of operators to perform the turnaround who did not have the required skills and expertise, inadequate use of lessons learned from the past, and audit procedures that did not identify the shortcomings before the accident.
- Improve the process control system to provide appropriate assistance to operators performing functions that are outside of normal production.

Shell Global (Corporate)

Three basic functions are included here: Engineering design (Shell Projects and Technology), corporate safety management, and executive-level corporate management, including the Board of Directors. The exact distribution of the safety responsibilities in the Shell Global management structure was not included in the accident report, so they may be distributed throughout the Shell Global management structure differently than assumed here. The bottom line is that they need to be somewhere.

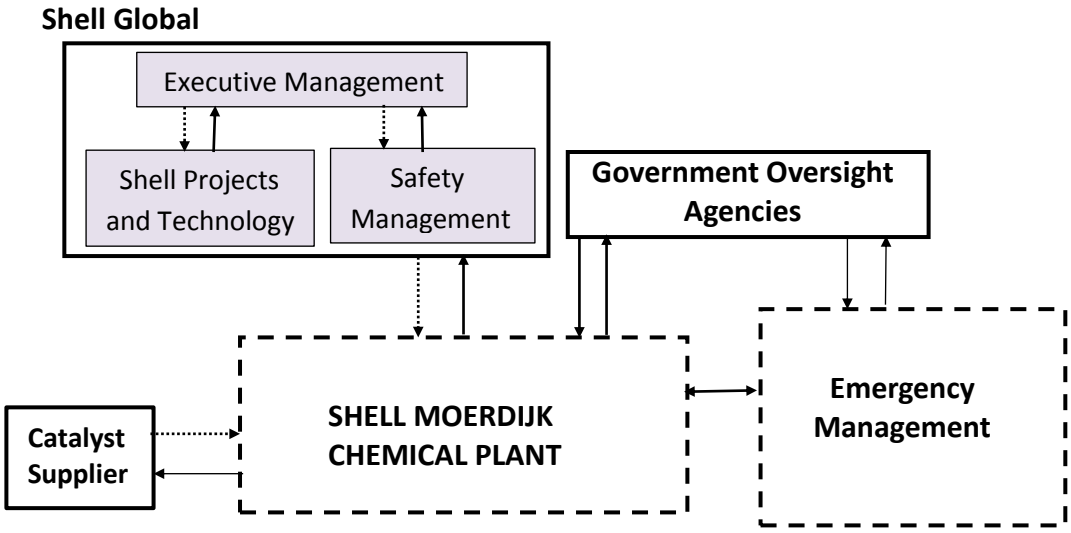

Shell Projects and Technology (Engineering)

Plant design was done at the corporate level and the technology licensed to the facilities.

Safety-Related Responsibilities

- Create a safe design: Perform hazard analysis (or use the results of hazard analysis created by another group) and eliminate or mitigate the hazards in the design.
- Provide design, hazard, and operating information to the plant operators to help those who are operating the plants avoid any hazardous scenarios that the designers were not able to eliminate or adequately mitigate in the design itself.
- Learn from the operation of their designs and improve the designs based on this feedback.

Summary of the Role of Shell Projects and Technology:

The design data provided to the licensees was not usable by those creating work instructions at the plants. The design had safety-critical design flaws that were not found in hazard analyses during the initial design phase. These flaws were not fixed after receiving information about related incidents in other Shell plants. One example was an overpressure incident that occurred due to an inadequate number of temperature sensors and pressure relief valves unable to handle the excessive pressure that occurred. Unsafe and incomplete work instructions were approved by Shell Projects and Technology for the Unit 4800 turnaround at Shell Moerdijk.

Unanswered Questions:

Without more information about the operations at Shell Corporate, which was not included in the accident report, it is difficult to determine exactly why the unsafe control occurred. More questions than answers arise from the CAST analysis, such as *Why were the design flaws introduced and how did they get through the design process? What type of hazard analysis is performed by Shell Projects and Technology or by other groups? Why were identified design flaws not fixed after the incidents at Shell Moerdijk in 1999 and Nanhai in 2011? What other types of feedback is provided about the safety of their designs during operations in the Shell plants? What information about the safety aspects (hazards) of the plant design are passed from Shell Projects and Technology to the licensees of their designs? What information is included in the design book? Is the design data provided sufficient for the licensees to create safe work instructions if engineers are writing the work instructions instead of operators and did they not know who was going to be performing this task? Why did they approve unsafe work instructions that did not even follow the required Shell format? What information is provided in the Design Book about start-up and the hazards of start-up? What types of hazard analyses are performed during the design process? What is the process for ensuring safety when changes are made? How are safety-related assumptions recorded and what triggers a re-analysis of these assumptions? What feedback do the designers get about the operation of their designs?*

Recommendations:

Fix the design features contributing to accident. Determine how these flaws got through the design process and improve both the design and the design review processes. Fix the design book so that is understandable by those who are writing the work instructions and includes all the information needed to safely operate installations of the licensed technology. Fix the work instruction review process by Shell Projects and Technology to ensure the instructions are complete and safe. Review and improve the hazard analysis process used by Shell Projects and Technology.

Corporate Safety Management

There is no mention in the accident report about a Shell corporate safety program or about any of its potential contributions to the accident. In the CAST analysis, I took the information about what happened at the local Shell Moerdijk level and projected what would normally be the responsibility at the corporate level of a well-designed SMS to control the flawed safety activities. There must have been someone with ultimate responsibility for safety at the corporate level, but it is unclear where the activities associated with that management role resided within the corporate structure or even whether these activities occurred.

Relevant Responsibilities

- Ensuring the safety of the plant design, including the conduct of hazard analysis on designs licensed to subsidiaries.
- Oversight of operational safety at the various Shell plants and facilities.
- Management of change procedures related to safety: creating them, making sure they are followed, and improving them using feedback from incidents.
- Ensuring communication among separate plants in different parts of the world about incidents, lessons learned, etc.
- Creating and updating a Shell-wide Safety Information System and ensuring the information is being communicated adequately both within Shell Corporate and globally and that it is complete and usable.

Summary of the Role of Corporate Safety Management: There appears to have been a flawed view of the state of risk and the effectiveness of the safety management system in local Shell plants. The flawed process model is most likely related to inadequate feedback (including audits and leading indicators). The accident report says that Shell uses the outdated HEMP and bow tie model. There is little information about what other hazard analyses and risk assessments are used at the corporate level. Presumably HAZOP is also practiced (as in most of the process industry), but that is not mentioned in the report. Bow tie, which is about 60 years old and dates back to the late 60s, uses a simple chain-of-events model and is too simplistic to capture the hazards and risks in today's complex systems, including chemical plants. Once again, many questions are raised from the CAST analysis that need to be answered to understand the role of corporate level safety management in the accident and thereby to provide more effective safety management in the future.

Recommendations: Improve Shell safety audits. Review all risk assessment and hazard analysis processes and, in general, improve their approach to safety with respect to both safety analysis and safety management. Shell is not alone among the large oil companies in needing to update their methods. The petrochemical industry has too many accidents and incidents that are avoidable.

More specifically, the accident report says that Shell should "evaluate how risk analyses are performed and make changes. This should include procedures and policies about re-evaluation of earlier presumptions and assumptions. Conduct new risk analyses, put adequate measures in place, and ensure that the team that performs these analyses has sufficient critical ability. Pay particular attention to assumptions based on risks that had previously been ruled out."

Evaluate and improve the corporate safety management system. Improve procedures for learning from process safety-related incidents. Create better feedback mechanisms, including audits and leading indicators, and procedures for learning from incidents.

Executive-Level Corporate Management

Responsibilities:

- Take all measures necessary to
  - Prevent major accidents from occurring and,
  - If accidents do occur, mitigate their consequences for humans and the environment
- Ensure the health and safety of the employees in relation to all aspects associated with the work (based on the Working Conditions Act and other regulations)
- Follow the government regulations in the countries where their plants are located.
- Create an effective safety management system and establish a strong safety culture policy. Ensure that the SMS and safety policies are being followed and they are effective.

Process Model Flaws

Leaders clearly had misunderstandings about the state of safety being practiced in Shell corporate and the local installations and the effectiveness of their defined procedures.

Contextual Factors (omitted, see full analysis):

Summary of the Role of Executive-Level Management:

Corporate management is responsible to ensure that an effective safety management system is created. Typical policies of an effective safety management system were violated at both Shell Corporate and Shell Moerdijk. The group overseeing safety at the Shell corporate level was clearly not effective. There is nothing included in the accident report about the assigned safety-related responsibilities for corporate management.

There is also nothing included in the accident report about context that might explain why standard executive-level responsibilities for safety were not effectively carried out. There seems to be a safety culture problem at Shell. (See later analysis of the safety culture and high-level safety policy at Shell.) What is the culture of the chemical industry in terms of corporate management oversight of the safety of global installations?

The accident report notes that the Safety Management System was integrated with the Business Management System at Shell Moerdijk. Was this also true at the corporate level? This is a very poor practice (and was a factor in the Deepwater Horizon accident). Safety risk assessments need to be kept separate from business risk assessments so that information is not hidden from high-level decision-makers.

Recommendations: Review the SMS design and determine why it did not prevent obvious violations of policy such as shortcomings in safety studies, management of change procedures, learning from accidents, not following regulations (e.g., having experienced operators and following the format for work instructions). Determine why audits were not effective in finding such obvious procedural noncompliance. While it is possible that this was the first time such lapses have occurred, it is highly unlikely. Strengthen audit procedures, including identifying better leading indicators of increasing risk than simply the number of leaks and create other forms of feedback to identify when the safety management system is drifting off course and risk is increasing. Establish better feedback channels to ensure that management of change procedures and corporate safety policy are being followed.

Catalyst Supplier

Safety-Related Responsibilities

- Provide information to customers that is needed to evaluate the use of their catalyst in the reactor being designed and/or operated
- Alert customers when changes are made in the catalyst that could potentially affect the safety of its use.

Summary of the Role of the Catalyst Manufacturer in the Accident: The changes made in the catalyst were not pointed out to Shell, but they were included in a new safety information sheet. While the catalyst manufacturer cannot determine the impact of their changes on a customer, the manufacturer should provide some clear alert that changes have been made and what they are so the customers are made aware of them.

Recommendations: Change contractual relationships between Shell and its suppliers to ensure that potentially critical changes are communicated appropriately. Make changes within information sheets so they are clear and obvious.

Dutch Regulatory Authorities

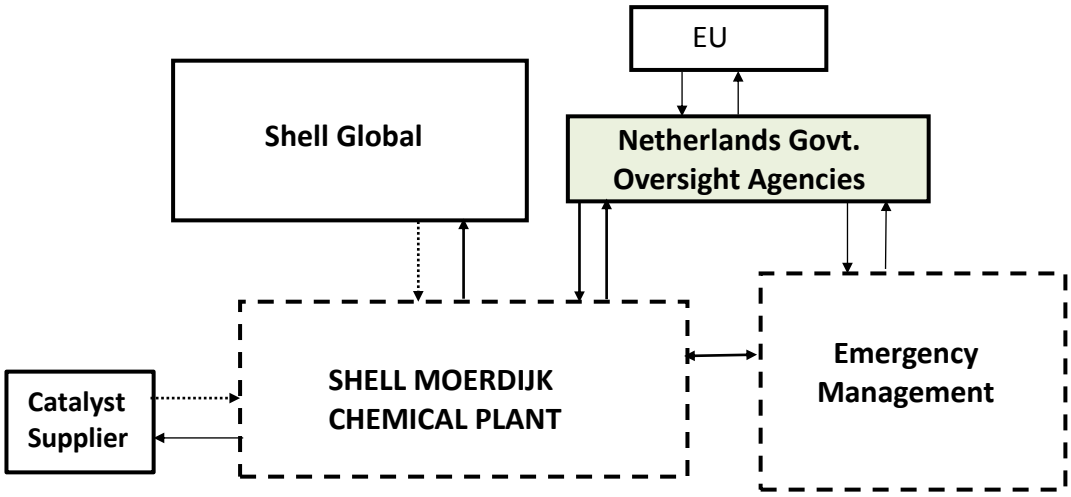

All Dutch oversight safety and environmental authorities are grouped together here. There are two main policies:

1. Brzo: Companies must take all measures to prevent accidents and, if they occur, mitigate their consequences for humans and the environment. The company must implement this obligation

by laying down policy assumptions in the Prevention Policy for Serious Accidents (PBZO), drawing up a safety report (VR), and organizing a safety management system.

2. Wabo: Regulators must check whether the company complies with regulations connected to the environmental permit, i.e., environmental safety.

General Relevant Safety-Related Responsibilities

- Responsible for supervision and enforcement of Dutch laws to protect the environment and the public. Perform Brzo inspections focusing on process safety and Wabo inspections focusing on environmental safety.
- Responsible for enforcement of EU health and safety laws within the Netherlands.

More Specific Responsibilities:

- Identify shortcomings at companies they are responsible to oversee.
- Encourage companies to improve their safety-critical processes through supervision and enforcement. Identify shortcomings and require companies to investigate and detect deep-seated causes of incidents and accidents. Ensure that any shortcomings identified are corrected.
- Assess modifications made to plants, procedures, and processes (although they are not expected to perform the risk analyses for the companies).
- Pay greatest attention to safety-critical processes, including maintenance and reactor start-up.

Process Model Flaws

The accident report said "Regulators had a positive view of the Shell Moerdijk safety management system. A number of shortcomings at Shell Moerdijk did not alter this view."

Contextual Factors (omitted, see full analysis and summary below)

Summary:

At least partly because of limited resources, the government authorities do "system-related supervision," effectively a form of performance-based regulation where responsibility is placed on the operator of high-risk activities to identify their own shortcomings. Regulators check both the design and operation of the safety management system and perform annual inspections to ensure they are operating as designed. Regulators only ensure that companies have the right documented procedures and spot check that they are being used.

They did not notice or did not react to Shell not acting in accordance with its own SMS. As just some examples: Changes and upgrades to the plant were not consistently subjected to risk analyses (violating the Shell SMS requirements), but this deficiency was not noted by the regulators nor required to be fixed. Changes were not adequately evaluated for safety. Requirements for expertise and training in performing startups were not enforced.

The accident report implies that regulators gave Shell Moerdijk a pass on behavior that might have instead been labeled violations. Plant scenario deficiencies should have been considered a violation but were not. Scenarios were not up to date or were incomplete. Working under limited resources and time is difficult under any supervision model, but system-level supervision has major limitations in ensuring public safety. The accident investigation showed many flaws in Shell Moerdijk operations

safety management as defined and as implemented. *What is wrong with the supervision model that the regulators did not detect the deficiencies?*

Recommendations:

Better supervision of the highest risk activities is needed, including turnarounds. Regulators need to oversee and ensure that strict procedures are being used for the most dangerous activities and that the safety management system is operating effectively and following its own rules. Operating under limited resources does not preclude doing something effective, it simply requires a more intelligent selection of activities that are performed. There is a need for better evaluation procedures and oversight of safety management system effectiveness. The regulators should rethink system-level supervision to ensure that what they are doing is effective in preventing accidents like the Shell Moerdijk explosion.

Emergency Services

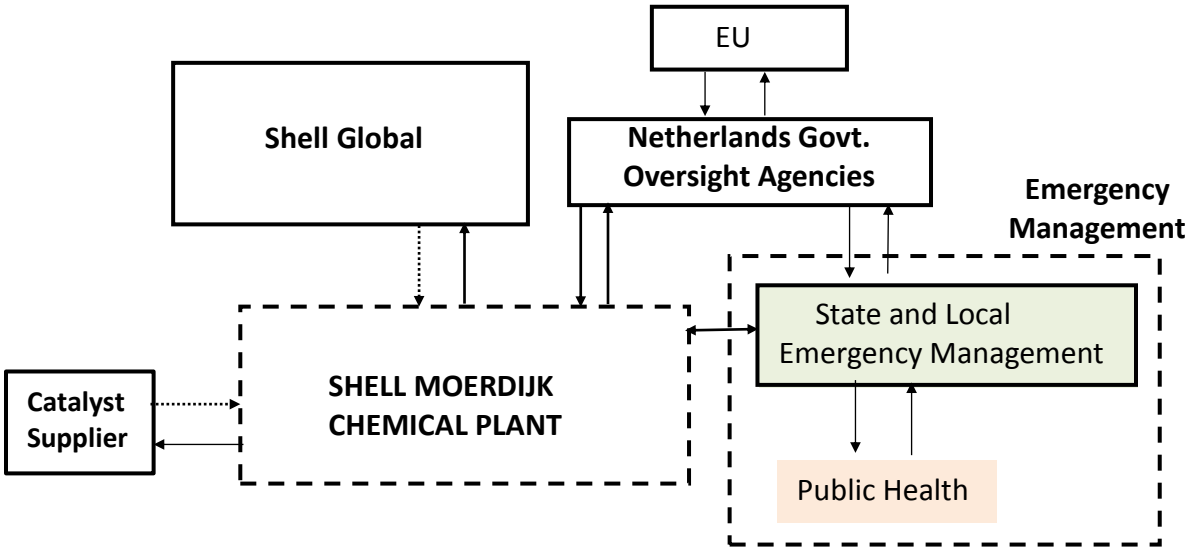

Responsibilities

- Firefighting [collaborative fire brigades did this effectively in this case], crisis management, crisis communications including among other things:
  - Informing citizens of the incident
  - Measuring substances released on a coordinated basis
  - Operating a telephone advisory line
  - Informing citizens about the results of the measurement of the substances released and the ensuing recommendations.

Summary of the Role of Emergency Services in the Accident:
Emergency services were mostly very effective in carrying out their responsibilities, but some deficiencies, particularly in communication, were uncovered in the accident response. For example, many people used WhatsApp instead of the National Crisis Management System designed for use in these conditions. It did not lead to loss of life because of the nature of the accident in this case, but under other circumstances it could have.

Recommendations: What was learned from this case should be used to improve the National Crisis Management System, including why many people used WhatsApp instead, and how the official system can incorporate those features.

Even for system components that did not cause a loss, as in this instance, accidents create an opportunity to look closely at the actual operation of our systems in times of stress and provide a learning opportunity that should not be wasted. As demonstrated in the example, most of the components may have contributed to the accident in some way or did not operate exactly as planned. When looking only at what the operators did or did not do, it appears that they were the primary contributors to the loss. Usually, as in this case, the actions of the operators will be found to be understandable when the reason for their behavior is examined.

This section described and provided examples of individual component analysis in CAST. Examining the individual controllers in the control structure is not enough, however. It is necessary to also examine the operation of the control structure as a whole.

Analyzing the Control Structure as a Whole

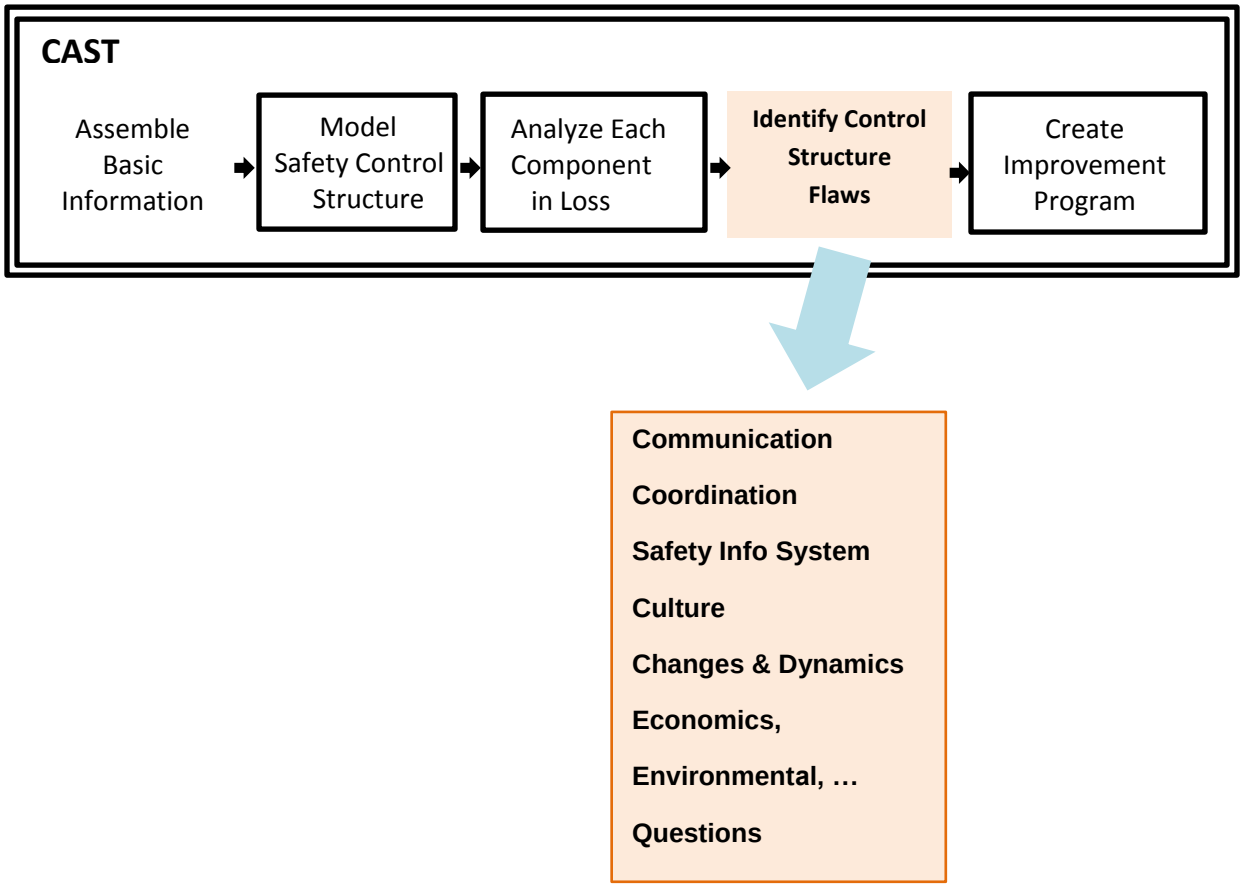

4. *Identify flaws in the control structure as a whole by investigating the general systemic factors that contributed to the loss. The systemic factors span the individual system control structure components.*

The CAST analysis process described so far focuses on individual system components and their control over other components. It examines why each of the controllers was unable to enforce the controls and constraints assigned to it.

In contrast, this part of CAST looks at the control structure as a whole and the systemic factors that led to the ineffectiveness of the designed controls. It is perhaps the least structured part of the analysis, but guidance is provided in this section of the handbook on what to look for and how to explain the role it played when you find it.

The systemic analysis focuses on factors that affect the behavior and interactions of all the components working together within the safety control structure to prevent hazardous system states. By looking at the system as a whole, rather than individual components, we can identify causal factors that impact how the various safety control structure components interact. These systemic factors provide another way to understand why the individual components may not fulfill their individual safety responsibilities and why together their behavior did not satisfy the system safety constraints. That is, safety control structures are usually created so that one component's misbehavior cannot alone lead to an accident. The systemic causal factors, however, can negatively impact the behavior of many or even all of the components and defeat these efforts.

This is the truly unique part of a systems approach to accident analysis that is omitted from event-based models, such as the Swiss Cheese or domino models. There are causal factors that can prevent all the barriers from operating correctly and simultaneously cause "holes" in all the protections and cheese slices that were created to prevent accidents. Getting away from focusing on a few causal factors or explaining accidents in terms of the behavior of one or several components provides a larger view of causality that can be used to identify very powerful prevention measures.

The following are some of the systemic factors that might be considered. This list is not comprehensive, but simply a starting point.

- Communication and coordination
- The safety information system
- Safety culture
- Design of the safety management system
- Changes and dynamics over time: in the system and in the environment
- Internal and external economic and related factors in the system environment not covered previously in the analysis. Some of these may be generated while considering changes over time.

Communication and Coordination

Many of the factors that lead to losses involve inadequate communication and coordination among components of the safety control structure. Lack of coordination and communication can lead to inconsistent actions, missing actions, and so on.

An example can be found in the collision of two aircraft in 2002 over Überlingen, a town in southern Germany in 2003. A poorly planned maintenance activity resulted in almost all the communication links in the air traffic control tower being temporarily broken. Some of the outages were known ahead, others resulted from an error during the maintenance. Planning for the loss of communication would have helped. For example, the air traffic controller on duty did not know that his aural and visual alerts about a potential collision would be inoperable during the maintenance. Contingency actions could have been taken to cope with the maintenance activity if it had been planned and management of change procedures used. The communication outage eliminated all the built-in redundancy that was designed to prevent such collisions.

Drawing the communication links as intended and what was operational during the accident can provide an important visual explanation. Figure 16 shows the communication channels that were designed into the system. Figure 17 shows those that were working at the time of the collision. This type of diagram is a powerful tool for understanding and communicating about what happened and why. The phone lines were out so that a controller at another site who saw the collision about to happen, could not get through to give a warning. When the air traffic controller who was controlling the two aircraft at the time of the collision became overloaded due to an unforeseen emergency landing, there was no way for him to contact another air traffic control facility to help him with his overload.

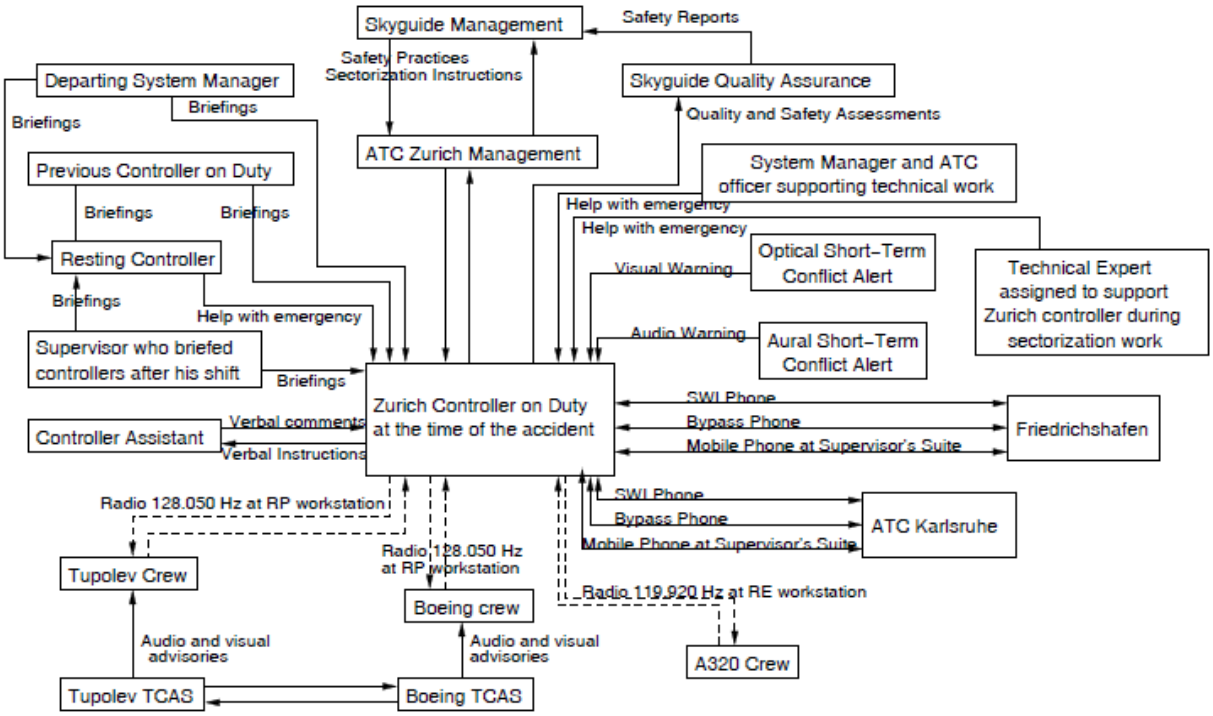

Figure 16: Communications links theoretically in place in the Überlingen accident

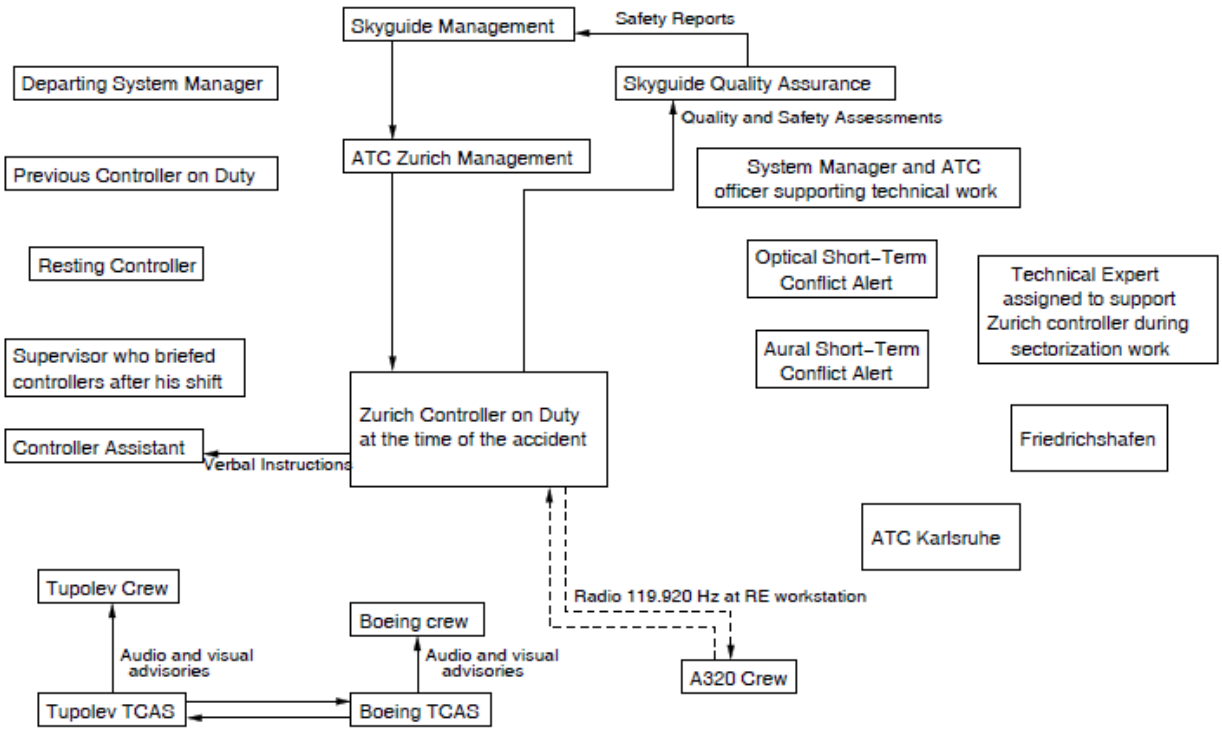

Figure 17: The operational communication links at the time of the accident

The communication and coordination problems in this case were confined to the time of the accident. But they can exist for a long time before. Sometimes, one controller assumes that another is handling a responsibility while the other one thinks the first is. As an example, in a Canadian water contamination incident, both the Ministry of the Environment (MOE) and the Ministry of Health (MOH) were responsible for performing some of the same oversight duties: the local MOH facility assumed that the MOE was performing this function and cut back on their own activities, but the MOE's budget had been cut, and follow-ups were not done. A common finding after an accident is that each of the responsible groups may assume another controller is performing the needed oversight when there are overlapping responsibilities. Alternatively, the two controllers may have non-overlapping responsibilities but they may provide indirectly conflicting commands.

Inadequate feedback in the safety control structure is one important type of communication that is often found to contribute to poor decision making. Other types of communication are, of course, also important. In the Shell Moerdijk accident, there were instances of ineffective communication between Shell Global Projects and Technology and Shell Moerdijk and between groups within Shell Moerdijk itself. In addition, the change to the catalyst by the manufacturer was not communicated to Shell. Did the required communication channels exist at the time of the accident? Were they fully operational or were there inhibitors to the transmission of required information?

In a friendly fire accident analyzed in Chapter 5 of *Engineering a Safer World*,[17] potential ambiguity in the responsibility for tracking friendly aircraft had been eliminated by partitioning the tracking responsibility between two different controllers. This partitioning broke down over time due to factors such as unofficial attempts to increase efficiency, with the result that neither was tracking the aircraft

[17] Nancy Leveson, *Engineering a Safer World*, MIT Press, 2012, Chapter 5.

that was shot down on the day of the loss. No performance auditing had been done to ensure that the assumed and designed behavior of the safety control structure components was actually occurring.

After an accident, all communication links should be examined to ensure that the design allowed for the transmission of critical information. Even if the design was adequate, were the transmission channels working properly? If not, the reasons for the inadequate communication must be determined. Sometimes a communication channel involves a person who has conflicting goals with the communication of the information. In some cases, the communication links may simply have failed. In others, they may be purposely down for maintenance purposes, as in the Überlingen accident. Occasionally, communication and coordination links may have been available, but the controller did not know about them. In the friendly fire example, the links degraded and became inactive over time.

Analyzing communication requires first understanding and modeling the critical communication channels. The control structure below, created by Captain Paul Nelson in his CAST analysis of a wrong runway takeoff accident at Lexington Airport, shows dotted lines where the communication channels never existed (some involving feedback channels that were never included in the design) or were inoperable for various reasons at the time of the accident. Simply drawing the safety control structure and finding missing feedback channels in the design is a strong clue that there is a problem with the control structure design. All control loops should have feedback channels for proper decision making to occur.

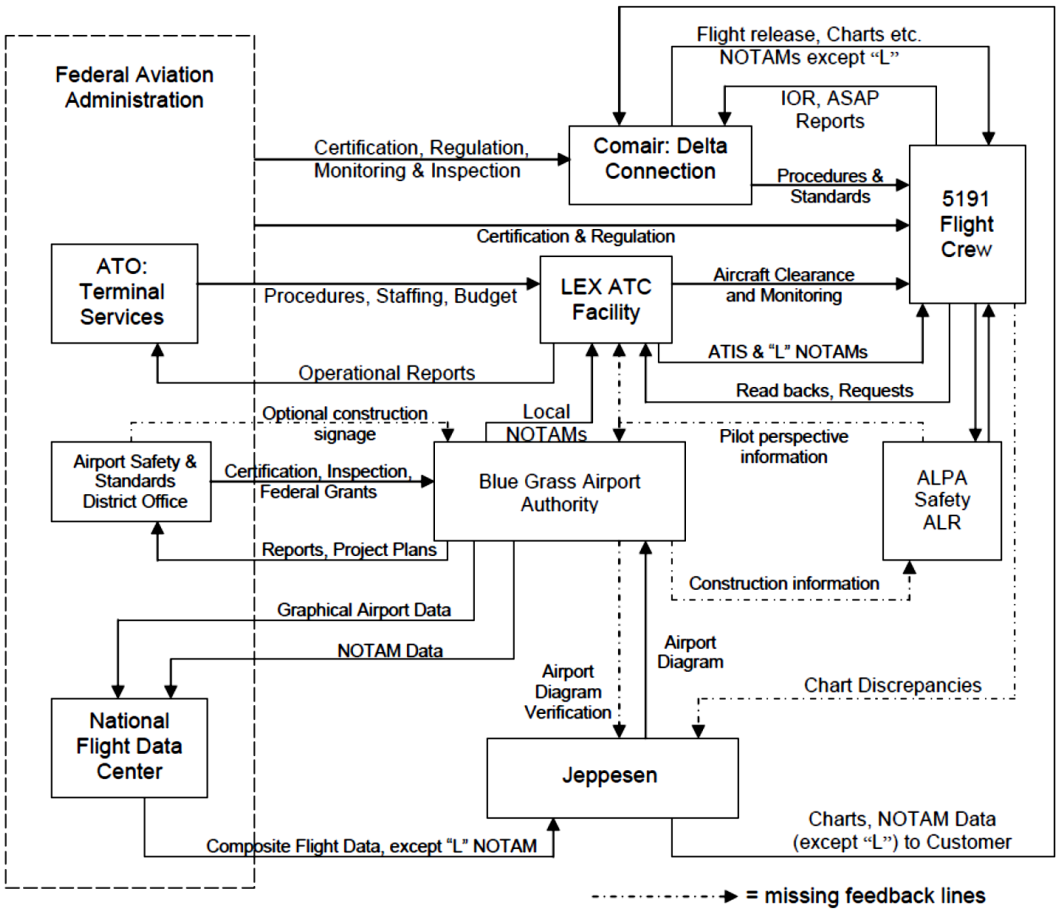

Figure 18: The Lexington ComAir wrong runway accident safety control structure. The control is drawn left to right instead of top to bottom. Dotted lines represent missing feedback and communication that contributed to the loss [Paul Nelson, A STAMP Analysis of the LEX Comair 5191

Accident, Master's Thesis, Lund University, June 2008. A link to this thesis is provided in Appendix A.]

Another common type of communication breakdown, described earlier, can occur in the problem-reporting channels. In many accidents, the investigators find that the problems were identified in time to prevent the loss, but the required problem-reporting channels were not used, often because they were unusable.

Safety information system

In a study of the safety information systems of various companies, Kjellan found that the quality of the safety information system was the second most important factor in discriminating between companies with high and low accident rates [Kjellan 1982].[18] Uses for a safety information system include storing and passing on information about hazards, detecting trends and deviations that portend an accident, evaluating the effectiveness of safety controls and standards, comparing models and risk assessments with actual behavior, identifying and controlling hazards to improve designs and standards, etc. Accident investigations should include an examination of the organization's safety information system (SIS).

The Shell Moerdijk accident report does not provide any information about the Shell or Shell Moerdijk SIS, but it does describe many instances of deficiencies that could have been prevented with a well-designed SIS, such as not learning from incidents and previous maintenance stops, not identifying flaws in previous hazard and risk assessments when contrary evidence arose, etc. The accident report also notes that important information was lost between the design of the unit and the ultimate operations management of the unit. The report says that "A discrepancy therefore occurred between the available information during the design phase and the operations management that was ultimately conducted." A well-designed safety information system should be able to prevent these types of problems.

Just having a safety information system does not guarantee that it will be useful or used. Data may be distorted by the way it is collected. Common problems are filtering, suppression, and unreliability. Data collected in accident and incident reports tend to focus on proximal events and actors, and not on the systemic factors involved such as management problems or organizational deficiencies. Limited categories of conditions may be identified, especially when checklists are used. It is difficult to obtain comparable data from multiple sources in an unbiased and systematic manner.

Data collection tends to be more useful for events that are similar to those that have occurred in the past than for events in new types of systems where past experience about hazards and causal factors is more limited. Software errors and computer problems are often omitted or inadequately described in incident reports because of lack of knowledge, lack of accepted and consistent categorizations for such errors, or simply discounting them as a causal factor. CAST can be used to encourage the inclusion of systemic factors in safety information systems.

Some common deficiencies include not recording the information necessary to detect trends, changes and other precursors to an accident; to evaluate the effectiveness of the controls used to prevent accidents; to compare risk assessments of those in the industry with actual behavior; and to learn from events and improve their safety management practices.

Simply collecting the information, of course, is not enough. Problems may arise from the difficulty in consolidating a large mass of data into a form useful for learning. While with digital technology today it

[18] The highest-ranking factor was top-level management concern about safety. Urban Kjellan, An evaluation of safety information systems at six medium-sized and large firms, *Journal of Occupational Accidents*, 3:273-288, 1982.

is easy to design large scale data collection channels, finding the time and manpower to analyze all the data that results may be difficult or impractical. As a result, the safety information system may contain only summary statistical data that can be easily processed by a computer but not the information about trends and changes over time that is needed to learn from events before major losses occur. Airlines today are particularly suffering from this type of data overload due to large-scale automated information collection.

Another deficiency of many safety information systems lies in the retrieval and dissemination mechanisms. Information—which is not the same as data—may not be presented in a form that people can learn from, apply to their daily operations, and use throughout the system life cycle. Updating may not occur in a timely manner: Accidents have resulted from changes during operations or due to insufficient updates to the hazard analyses when engineering modifications were made. The information may not be tailored to the needs and cognitive styles of the users or not integrated into the environment in which safety-related decisions are made and therefore may be hard to use.

The safety information system must also exist within a safety management system that can use and benefit from the information provided. Simply having information stored in a safety information system does not mean that companies have the structures and processes needed to benefit from it.

When conducting an accident investigation, the safety information system and any possible impact on the events needs to be examined. Was the information needed to prevent the loss not collected or stored? Was it lost in the collection or analysis process? Was it available and easily retrieved in the daily activities of controllers?

Design of the safety management system

The safety management system (SMS) is theoretically the same as the safety control structure used in CAST analyses. The more general term "safety control structure" is used here as some industries define an SMS that excludes important controls necessary to prevent accidents.

There is no single correct design for the safety control structure: Alternative safety control structures can be effective with responsibilities distributed in different ways. The culture of the industry and the organization will play a role in what is practical and effective. There are some general design principles, however, that are necessary for any safety control structure to be effective,[19] and these can be used to evaluate an organization's safety management system after an accident. In general, the necessary safety constraints on organizational behavior should be reflected in the design of the safety control structure: Is the assignment of responsibility, authority, and accountability for safety to the controllers adequate to prevent hazards and accidents? A list of general responsibilities for management, development and operations that need to be assigned is included in Appendix D.

The accident investigation should evaluate the official design of the safety control structure as well as how it actually operated at the time of the loss. Were performance or other types of audits used to ensure that it was working effectively? Previous minor incidents may be clues that something is amiss and should have been used to identify problems before a major loss. Were near-miss and incident reporting systems used to identify flaws in the operation of the safety control structure or simply to assign blame to operators?

Clues will be available during the investigation that will provide guidance on where to look. For example, the Shell Moerdijk official accident report, which did not document or evaluate Shell's safety management system, notes that unsafe situations were overlooked, internal procedures were not properly followed, lessons were not learned from previous incidents, incorrect assumptions about basic

[19] See Leveson, *Engineering a Safer World, 2012* and Chapter 7 (Designing and Effective Safety Management System) in Leveson and Thomas, *STPA Handbook*, 2018.

chemical reactions were not re-evaluated after evidence surfaced that they were incorrect, changes were not managed and controlled, inadequate hazard analysis and risk assessment procedures were used, recommendations from previous incidents and accidents were never implemented, and oversight of critical activities was missing. Were responsibilities assigned for performing and overseeing these activities? What types of audits or inspections were used to identify these types of poor practices? Why were they not identified and corrected before the accident?

Oversight of a company's safety management system (SMS) usually lies in the responsibilities of government or professional regulatory agencies or organizations. At Shell Moerdijk, for example, the Dutch Regulatory Authorities have the responsibility to check whether the company has a safety management system in place, whether the systems and procedures incorporated in that system are appropriate, and whether the company actually applies these systems and procedures. They did not notice or did not react to Shell not acting in accordance with its own SMS. As just some examples, changes and upgrades to the plant were not consistently subjected to risk analyses, which violated the Shell SMS requirements, but this deficiency was not noted by the regulators nor required to be fixed. Changes were not adequately evaluated for safety. Requirements for expertise and training in performing startups were not enforced.

In fact, the regulators were unanimous in their positive appraisal of the Shell Moerdijk SMS. Government inspections and supervision are partly determined by the judged quality of the SMS, so oversight activities were limited for Shell Moerdijk. The Dutch Regulatory Authorities did not have adequate resources (a common problem) and thus allocated them based on perceived risk. They only perform "system-oriented supervision," that is, if the right systems are in place, then the authorities do not look further at the actual activities in the plant. The official accident report notes, however, that even under this type of limited supervision, it was possible for the inspectors to observe that changes and upgrades to the plant were not consistently subjected to risk analyses and that the safety management system did not indeed function well.

In the Shell Moerdijk accident, there had been two previous incidents involving the same catalyst, neither of which seemed to generate concerns or questions about the design of the plant, the risk assessments that were performed, or the creation of the start-up work instructions. The information about the incidents may not have been effectively stored by Shell (and thus stemmed from a deficiency in the safety information system), but more likely it simply was not used due to deficiencies in the operation of the company's safety management system. Asking appropriate questions during the investigation could have identified the source of these problems.

Safety culture

Acceptable safety information systems and safety management system structures can both be defeated by safety culture problems. In fact, the safety culture in an organization or industry impacts the behavior of every component in the safety control structure.

Edgar Shein, considered to be the "father of organizational culture," defined safety culture as "*the values and assumptions in the industry and/or organization used to make safety-related decisions*"[20] Figure 19 shows Shein's three levels of organizational culture and, more specifically, safety culture.

[20] Edgar Shein, *Organizational Culture and Leadership*, San Francisco: Jossey Bass, 2004

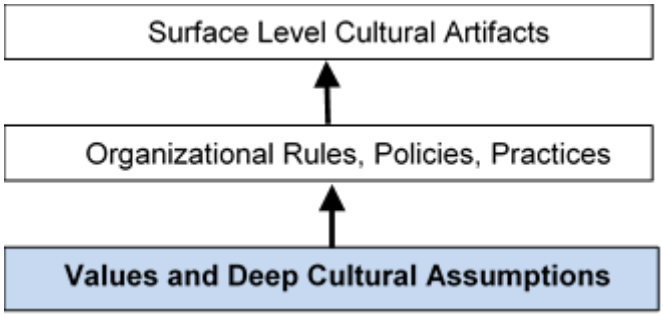

Figure 19: Shein's model of organizational culture

The essence of culture is the lowest-level set of values and deep cultural assumptions that underlie decision-making and the higher-level cultural artifacts in an organization or industry. The middle and top levels are not the culture; they merely reflect the organizational culture.

There is always a safety culture. The important question is how the existing safety culture impacts accident rates. Some aspects of safety culture that you might look for when investigating an accident are:

- Culture of Risk Acceptance: This culture is based on the assumptions that accidents are inevitable and that nothing much can be done to prevent them beyond exhorting everyone to be careful. Accidents are considered to be the price of productivity. Often this assumption is accompanied by the belief that everyone should be responsible for safety—their own and others—and that accidents result from a lack of responsible behavior on the part of individuals. The belief is prevalent that if only everyone would act responsibly and safely, accidents would be reduced or even eliminated.
- Culture of Denial: In a culture of denial, risk assessment is often unrealistically low, with credible risks and warnings being dismissed without appropriate investigation: Management only wants to hear good news so that is what they are told. The focus is on showing the system is acceptably safe, not on identifying the ways it might be unsafe. The use of "safety cases" rather than hazard analyses, is common.
- Culture of Compliance: The focus here is on complying with government regulations. The underlying cultural belief is that complying with regulations will lead to acceptable results. Because regulatory agencies tend to focus on certifying that a final product or service is safe, assurance of the safety of completed or existing designs is emphasized rather than building safety into the design during development. Often extensive "safety case" arguments are produced for managers or regulators with little or no impact on the actual product or product design effort. In addition, the bulk of the safety efforts may involve complying with standards and the requests of regulators instead of proactively taking steps to improve safety because it is a basic value in the organization. Even if employees fix immediately what is found to be lacking by government inspectors, did they aggressively search for problems internally without being prompted by a regulatory agency?
- Paperwork Culture: This culture rests on the belief that producing lots of documentation and analysis paperwork leads to safe products and services. Piles of paper analyses are produced but they have little real impact on design and operations. The bulk of the safety-related paperwork may be produced by a group that is independent from and has little interaction with those who

are designing and operating the products, implementing the processes, or providing the services. The paperwork sits in files, assuring everyone that the system is safe.

- Culture of "Swagger": Safety is for sissies. Real men thrive on risk. The problem here with respect to preventing accidents is, of course, obvious.

The safety culture in any organization is set by the top management. A sincere commitment by management to safety is often cited as the most important factor in achieving it. Employees need to feel that they will be supported if they exhibit a reasonable concern for safety in their work and if they put safety ahead of other goals such as schedule and cost. Often, concern for safety is merely lip-service and sloganeering. Employees take their cues from actions, not from empty words.

Analysis of the safety culture during an accident investigation should start with an evaluation of the company's documented safety philosophy and safety policy. If these do not exist, then there clearly is a problem. After a serious safety problem was found in a medical device, I was asked by the company to come in and investigate what happened. I had told them previously that they needed a written safety policy, but they had never produced one. After the safety problems in their products, I asked why and was told that it should have been obvious to the employees that the devices they were designing and selling were safety-critical and therefore a policy was not needed. At the same time, there were lots of written policies and rules about other product properties. In addition, they said that their lawyers had advised them not to create a safety policy because it was an admission that their product was potentially hazardous. The hazardous nature of the product was obvious and the lack of a safety policy would not have fooled anyone into thinking that there were no hazards associated with it. It simply made the company look negligent. Shortly after the adverse events that I was consulted about, the company went bankrupt and ceased to exist.

The safety philosophy is a short, written statement of the relationship that management desires between safety and other organizational goals. Some examples of general principles that might be part of the Safety Philosophy statement:

1. All injuries and accidents are preventable.
2. Safety and productivity go hand in hand. Improving safety management leads to improving other quality and performance factors. Maximum business performance requires safety.
3. Safety has to be built into a product or the design of a service. Adding it later will be less effective and more expensive. After-the-fact assurance cannot guarantee a safe design where safety is not already present. It is better to build safety in than try to ensure it after the fact.
4. The goal of accident/incident causality analysis is to determine why the loss (or near loss) occurred so that appropriate changes can be made rather than to find someone or something to blame.
5. Incidents and accidents are an important window into systems that are not operating safely and should trigger comprehensive causal analysis and improvement actions.
6. Safety information must be surfaced without fear. Safety analysis will be conducted without blame.
7. Safety commitment, openness and honesty is valued and rewarded in the organization
8. Effective communication and the sharing of information is essential to preventing losses.

While these principles look deceptively simple, they actually take quite a bit of thought to create and are often not the principles actually guiding a company's decision making. More explanation behind them can be found in the STPA Handbook, Chapter 7.

The safety philosophy establishes the principles upon which the more extensive safety policies and standards are built and assessed and provides a succinct statement of how management wants employees to make safety-related decisions and to act.

Of course, simply having a written philosophy does not mean that there has been buy-in by management and other employees. Do the managers actually implement the philosophy in their daily practices? Is there any process for ensuring that the philosophy is adopted and the principles practiced? How are the principles communicated to the employees? A written statement of the safety philosophy and more detailed policy statements and standards is a start, but it is not enough. Employees quickly identify when the written policy differs from the actual behavior of management. Do employees believe those at the top of the organization are sincerely committed to safety and are not just sloganeering and going through the motions? Asking some simple questions can identify when a deficient safety culture contributes to accidents.

While accident investigations sometimes use elaborate employee surveys to understand the existing safety culture, I have never found them useful. Writing down in a survey what one thinks they should say or believe does not mean that they actually think or act in that manner. I find that observing actual behavior is much more enlightening. Safety culture is best evaluated by studying management and employee behavior directly. How people behave is a much better indicator of their internal value system than what they answer on surveys.

How is commitment demonstrated? It is shown by setting priorities and following through on them; by personal involvement (e.g., top management chairing groups where safety decisions are made or at least getting personally involved); by setting up appropriate organizational structures; by appointing designated, high-ranking leaders to safety-related responsibilities and providing adequate resources for them to be effective; by assigning the best employees to safety-related activities and rewarding them for their efforts; and by responding to initiatives by others. It is also communicated by minimizing blame. Leaders need to demonstrate that their highest priority is to fix the systemic factors leading to losses and not just to find someone on which to pin blame (usually someone at the lowest levels in the organization) and then moving on when incidents or accidents occur. Finally, the incentive structure in the organization should be designed to encourage the behavior desired. Again, asking some questions about these factors during an accident investigation can provide insight into the contribution of the safety culture to the events. Does management value safety enough to use it as an important criterion when making decisions about promotions and raises?

Changes and Dynamics over Time: In the System and in the Environment

Accidents usually occur after some type of change. The changes may be in the physical process, the operating procedures, individual behavior, the safety activities or processes, the management process, oversight practices (both internal and external), or in the environment in which the system exists and with which it must interact.

Changes may be planned or unplanned. Both types of changes need to be controlled and can lead to accidents if they are not.

If the changes are planned, a strong and well-designed management of change policy that is enforced and followed should be in place. In many accidents, management of change (MOC) procedures existed, but they were neither effective nor enforced. Examples in the Shell Moerdijk explosion include the switch to a new catalyst without testing it and assuming that previous catalyst properties still held and the removal of parts of the work instructions for Unit 4800 (again without assessment) because they were not considered critical. Critical requirements regarding nitrogen flow were removed during periodic updates of the work instructions in an attempt to limit their content to information that was

believed essential and to focus on what was incorrectly thought to be the most important from a safety and operational view. Other information was omitted from the work instructions because, over time, understanding of the most appropriate procedures related to Unit 4800 changed.

Changes impacting risk may also be unplanned. There needs to be a way to detect unplanned changes that affect safety and to prevent them from occurring. Detection may be accomplished by using leading indicators and safety-focused audits. There may also be periodic planned re-evaluation of assumptions underlying the original safety-related design features and management procedures. At Shell Moerdijk, the leading indicators (such as number of leaks) were inadequate and too narrow, audits did not seem to be effective, and assumptions about the properties of ethylbenzene established in 1977 were never revisited.

Changes may occur slowly over time, as occurred at Shell Moerdijk with the work instructions for Unit 4800. As the work instructions were amended before each turnaround, important information was omitted—in some cases intentionally and in others unintentionally. Examples include the nitrogen flow requirements mentioned above and the required heating rate for the reactor. Changes do not appear to have been reviewed by experts, but if they were, then the review process was flawed.

Changes may be known and planned in one system component but appear as unplanned and unknown changes to another component of the system. The change in composition of the catalyst was known by the catalyst manufacturer but not by Shell Moerdijk. Clearly communication is an important factor here.

Leading indicators are commonly used in some industries to identify when the system is migrating toward a state of higher risk. At Shell Moerdijk, as well as most chemical plants, leading indicators used may be common throughout the whole industry, in this case number of leaks. Their selection seems to be predicated primarily on ease of collection and lack of having any alternatives for creating more effective ones.

I have created a new approach to identifying leading indicators that I called "assumption-based leading indicators."[21] Briefly the idea is that certain assumptions are made during system development that are used to design safety into a system. When, over time, those assumptions no longer hold, then the organization is likely to migrate to a state of higher risk. Leading indicators, then, can be identified by checking the original safety-related assumptions during operations to make sure that they are still true. Systems will always change and evolve over time, as will the environment in which the system operates. Because changes are necessary and inevitable, processes must be created to ensure that safety is not degrading. This new approach to leading indicators is being evaluated by Diogo Castilho, a Ph.D. candidate at MIT, on airline operations.[22]

Changes may evolve slowly over time and their impact may not be obvious. In a CAST analysis of crash of a cargo aircraft while landing at Birmingham Airport, one factor implicated in the inadequate operation of the safety control structure as a whole was an increase in night cargo operations at airports. Is there as much weight placed on cargo aircraft safety as passenger aircraft. Is more concern shown for daylight operations than for operations in the early darkness, which may be complicated by fatigue? We found that historical assumptions about airport operations may need to be revisited in the light of changes to airline operations and traffic.

Informal (and even formal) risk assessment may change as time passes without a loss. The actual risk has probably not decreased, but our perception of it does. That leads to a change in priorities and in the

---

[21] Nancy Leveson, (2015), A systems approach to risk management through leading safety indicators, *Reliability Engineering and System Safety*, 136: 17-34, April.

[22] Diogo Silva Castilho, *A Systems-Based Model and Processes for Integrated Safety Management Systems*, Ph.D. Dissertation, Aeronautics and Astronautics, MIT, expected September 2019.

resolution of conflicting goals. Indeed, the Space Shuttle losses can be partly explained in this way, particularly the Columbia accident. Unfortunately, a strange dynamic can arise where success at preventing accidents can actually lead to behavior and decision making that paradoxically increases risk. A circular dynamic occurs where safety efforts are successfully employed, the feeling grows that accidents cannot occur, which leads to reduction in the safety efforts, an accident, and then increased controls for a while until the system drifts back to an unsafe state and complacency again increases and so on.

The complacency factor is so common that safety control structures need to include ways to deal with it. SUBSAFE, the U.S. nuclear submarine safety program puts major emphasis on fighting this tendency and has been particularly successful in accomplishing this goal. Identifying the migration to states of higher risk is an important part of any accident investigation. Understanding the reason for this migration during an accident causal analysis in which unrealistic risk perception is involved can help to identify ways to design the safety control structure to prevent it or detect it when it occurs.

One way to combat this erosion of safety is to provide ways to maintain accurate risk assessments in the process models of the system controllers. The better information controllers have, the more accurate will be their process models and therefore their decisions. Accident analysis should involve a careful investigation of the accuracy of risk perception in the mental models of those controlling safety.

Internal and external economic and related factors

Systemic factors contributing to accidents often involve both internal and external economic conditions such as external market competition or declining profits. Market conditions may lead to management reducing the safety margins and ignoring established safety practices. Usually, there are precursors signaling the increasing risks associated with these changes, but too often these precursors are not recognized.

Another common factor is reduction in the physical separation between people and dangerous processes. Hazardous facilities are usually originally placed far from population centers, but the population shifts after the facility is created. People want to live near where they work and do not like long commutes. Land and housing may be cheaper near smelly, polluting plants. In third world countries, utilities, such as power and water, and transportation may be more readily available near heavy industrial plants, as was the case at Bhopal. The lure of providing jobs and economic development may encourage government officials to downplay risks and not rigorously enforce their safety control and emergency response requirements.

Over time, the land available for building may only be available near dangerous parts of the plant, as happened at the Texas City refinery where trailers were built to house employee offices next to the ISOM tower that exploded in 2005. With increasing population, local emergency facilities, such as firefighting and medical resources, may lag behind the increasing requirements due to constraints on resources and other competing priorities.

These factors may relate to the system environment not covered previously in the analysis of the individual controllers or even other general systemic factors, particularly changes and dynamics over time. How was individual controller behavior affected by these factors? Obvious factors are financial and competition concerns as well as new products and markets. Some non-obvious factors might be fatigue and duty cycles, distraction by contract negotiations, poor morale on the part of the work force due to a variety of factors, pressures to increase productivity beyond a reasonable level, etc.

Sometimes processes are automated and the standard safety controls implemented by humans are not implemented in the automated processes. There is often a naïve belief that software is always safe because it doesn't "fail" or that computers are more reliable than humans and therefore will improve

safety. Unfortunately, introducing automated systems introduces other changes into the system, often some that were not previously present. Complacency rises and new causes of accidents are introduced that are not controlled by the existing safety control structure.

Generating Recommendations and Changes to the Safety Control Structure

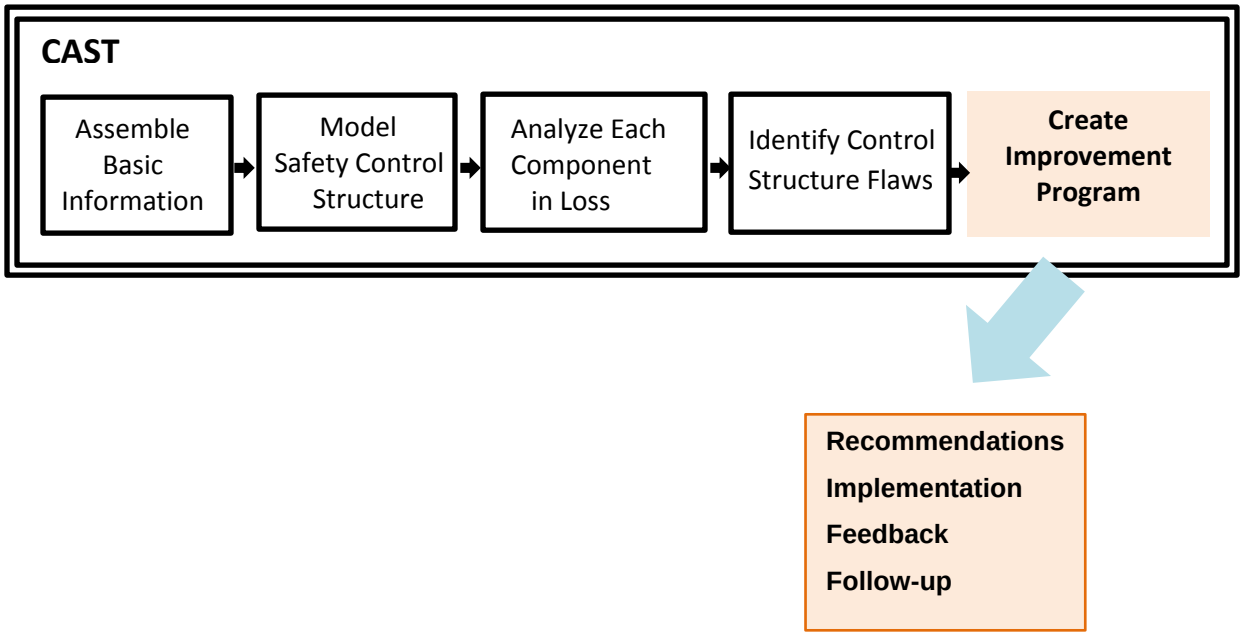

5. *Create recommendations for changes to the control structure to prevent a similar loss in the future. If appropriate, design a continuous improvement program for this hazard as part of your overall risk management program.*

Generating Recommendations

Once the other parts of the analysis are completed, generating recommendations should be straightforward. The biggest complaint about CAST we hear is that it generates too many recommendations. This complaint is only justified if the goal of the accident investigation is to make as few recommendations as possible. Accident investigation has had the goal of identifying "root causes" or "probable causes" for so long that it may be a cultural shock to start looking at more causal factors that generate more recommendations.

One of the objections raised to including a large number of recommendations is that responding to them is overwhelming. This is simply a logistical problem and not one that should be solved by learning less from each accident. There is no reason that recommendations cannot be prioritized according to specified criteria. There is also no implication that all the recommendations must be implemented immediately. Some recommendations will be relatively straightforward to implement immediately while others may take longer. Some may require such extensive changes that implementing them will take a great deal of effort and resources. Examples of the latter include establishing a new oversight agency or changing regulations and laws. Difficulty of implementation is not an excuse to omit a recommendation from the accident report, but it may be a good reason to categorize it as a longer-term goal rather than an immediate fix.

Taking steps to “jury-rig” short-term solutions should not be an excuse for endlessly delaying comprehensive and effective solutions.

Establishing a Structure for Continual Improvement

Sometimes recommendations are made but never implemented. Not only must there be some way to ensure recommendations are followed, there must also be feedback to ensure that they are effective in terms of achieving the goals and strengthening the safety control structure.

Essentially there are three requirements:

1. Assigning responsibility for implementing the recommendations
2. Checking that they have been implemented
3. Establishing a feedback system to determine whether they were effective in strengthening the controls.

The third requirement implies the need to collect evidence about the effectiveness of the recommended changes. Such feedback can come from audits and inspections and from the analysis of later incidents to determine whether previous recommendations were successful. Such an activity is a critical component of any safety management system, but it is often omitted.

Subsequent accidents or losses, particularly if they are analyzed using CAST, provide a rich source of information to understand why previous recommendations were not effective and what else is needed. Was the original causal analysis flawed? Were assumptions about the potential effectiveness of particular improvements incorrect? Did other changes occur that thwarted the attempt to strengthen the safety control structure? Did the planned changes result in unforeseen consequences?

The goal here is to ensure that your organization is continually learning and improving its risk management efforts so that the potential for losses is reduced over time.

Suggestions for Formatting the CAST Results

In STAMP, losses involve complex processes, not just a simple chain of events or even combination of event chains. The interactions among the events and causal factors are often intricate and subtle or indirect. In the end, it may be infeasible to list them as separate factors but only to understand the relationships among them. As a result, the format for presenting CAST results should emphasize the relationships among the various causal factors rather than a simple list, although lists are often difficult to avoid.

Tools for storing and organizing the results of a CAST analysis could easily be created, with extensive use of hyperlinks to show the relationships between various factors and behaviors and the location of answers to generated questions throughout the CAST analysis. More difficult, however, is coming up with a satisfactory standard notation for the presentation of the CAST results.

The goal is, of course, to provide the reader with an understanding of all the causal factors involved and how they all relate to each other. Different notations have had advantages in different accident analyses, with none that we have found working best for all of them.

As a result, documenting all the details of the analysis can use various types of notation, but readability and comprehensibility are important in a final summary of the causes of the accident. While all the details may not be necessary, there needs to be a way to present critical conclusions and argument. Trying to fit this all on a page (as with AcciMap analyses) leads to oversimplification and omission of important information. On the other hand, spreading the information over too many pages may lead to information overload and missing the bigger picture. I present here only a couple of examples that I have used and leave it for others to come up with better ideas.

For the Shell Moerdijk accident, I used the following format: a colored safety control structure with a table using colors as links to the summary analyses of the components as can be seen in Appendix B.

Alternatively, I have tried to insert the most important summary information in the control structure itself. Several figures from my analysis of a friendly fire accident where a U.S. F-15 shot down a U.S. Black Hawk helicopter over the Iraqi No-Fly-Zone are included here. The first shows the actions and analysis for the pilots, i.e., the F-15 lead pilot, the F-15 wing pilot, and the Black Hawk pilots. The second shows the analysis of the director of the ACE (the Airborne Command Element) up in the AWACS aircraft, who was commanding traffic in the area. The Mission Director was on the ground and overseeing the activities of the ACE. The third figure shows the analysis of the AWACS crew. Each of these figures is described in much more detail with accompanying text in the actual CAST analysis report. The notation does, however, allow an overview of the analysis. The additional text in the full CAST report should provide important but more detailed information.

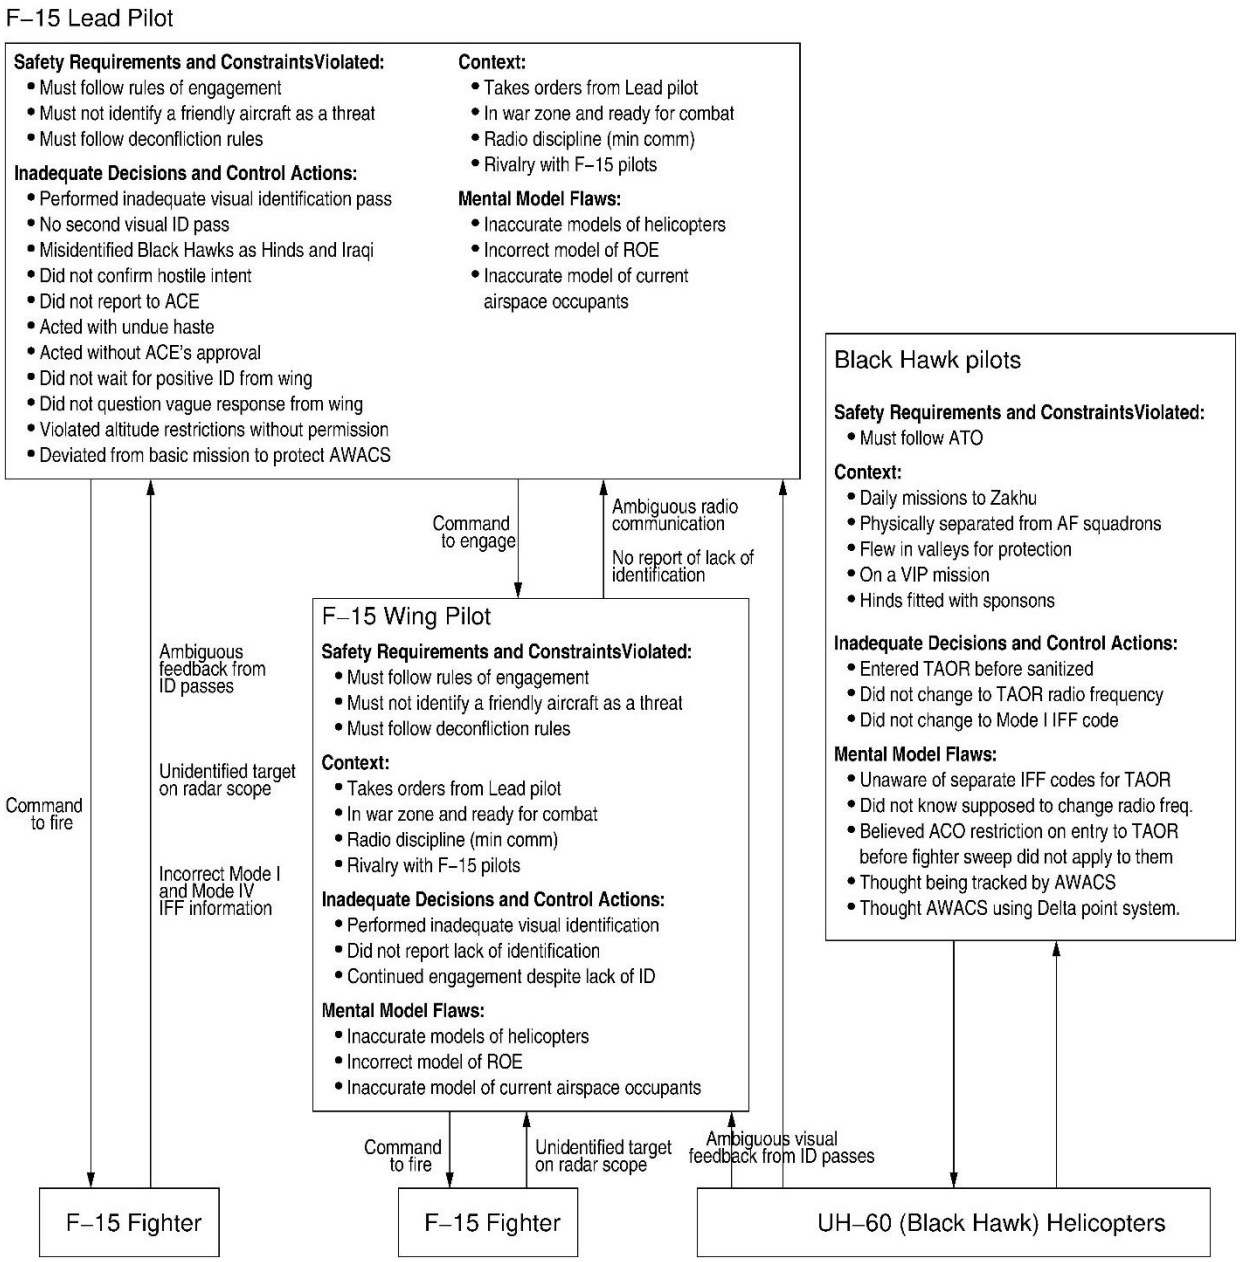
F–15 Lead Pilot
Safety Requirements and ConstraintsViolated:
• Must follow rules of engagement
• Must not identify a friendly aircraft as a threat
• Must follow deconfliction rules
Inadequate Decisions and Control Actions:
• Performed inadequate visual identification pass
• No second visual ID pass
• Misidentified Black Hawks as Hinds and Iraqi
• Did not confirm hostile intent
• Did not report to ACE
• Acted with undue haste
• Acted without ACE's approval
• Did not wait for positive ID from wing
• Did not question vague response from wing
• Violated altitude restrictions without permission
• Deviated from basic mission to protect AWACS
Context:
• Takes orders from Lead pilot
• In war zone and ready for combat
• Radio discipline (min comm)
• Rivalry with F–15 pilots
Mental Model Flaws:
• Inaccurate models of helicopters
• Incorrect model of ROE
• Inaccurate model of current airspace occupants
Black Hawk pilots
Safety Requirements and ConstraintsViolated:
• Must follow ATO
Context:
• Daily missions to Zakhu
• Physically separated from AF squadrons
• Flew in valleys for protection
• On a VIP mission
• Hinds fitted with sponsons
Inadequate Decisions and Control Actions:
• Entered TAOR before sanitized
• Did not change to TAOR radio frequency
• Did not change to Mode I IFF code
Mental Model Flaws:
• Unaware of separate IFF codes for TAOR
• Did not know supposed to change radio freq.
• Believed ACO restriction on entry to TAOR before fighter sweep did not apply to them
• Thought being tracked by AWACS
• Thought AWACS using Delta point system.
Command to engage
Ambiguous radio communication
No report of lack of identification
F–15 Wing Pilot
Safety Requirements and ConstraintsViolated:
• Must follow rules of engagement
• Must not identify a friendly aircraft as a threat
• Must follow deconfliction rules
Context:
• Takes orders from Lead pilot
• In war zone and ready for combat
• Radio discipline (min comm)
• Rivalry with F–15 pilots
Inadequate Decisions and Control Actions:
• Performed inadequate visual identification
• Did not report lack of identification
• Continued engagement despite lack of ID
Mental Model Flaws:
• Inaccurate models of helicopters
• Incorrect model of ROE
• Inaccurate model of current airspace occupants
Ambiguous feedback from ID passes
Unidentified target on radar scope
Incorrect Mode I and Mode IV IFF information
Command to fire
Command to fire
Unidentified target on radar scope
Ambiguous visual feedback from ID passes
F–15 Fighter
F–15 Fighter
UH–60 (Black Hawk) Helicopters

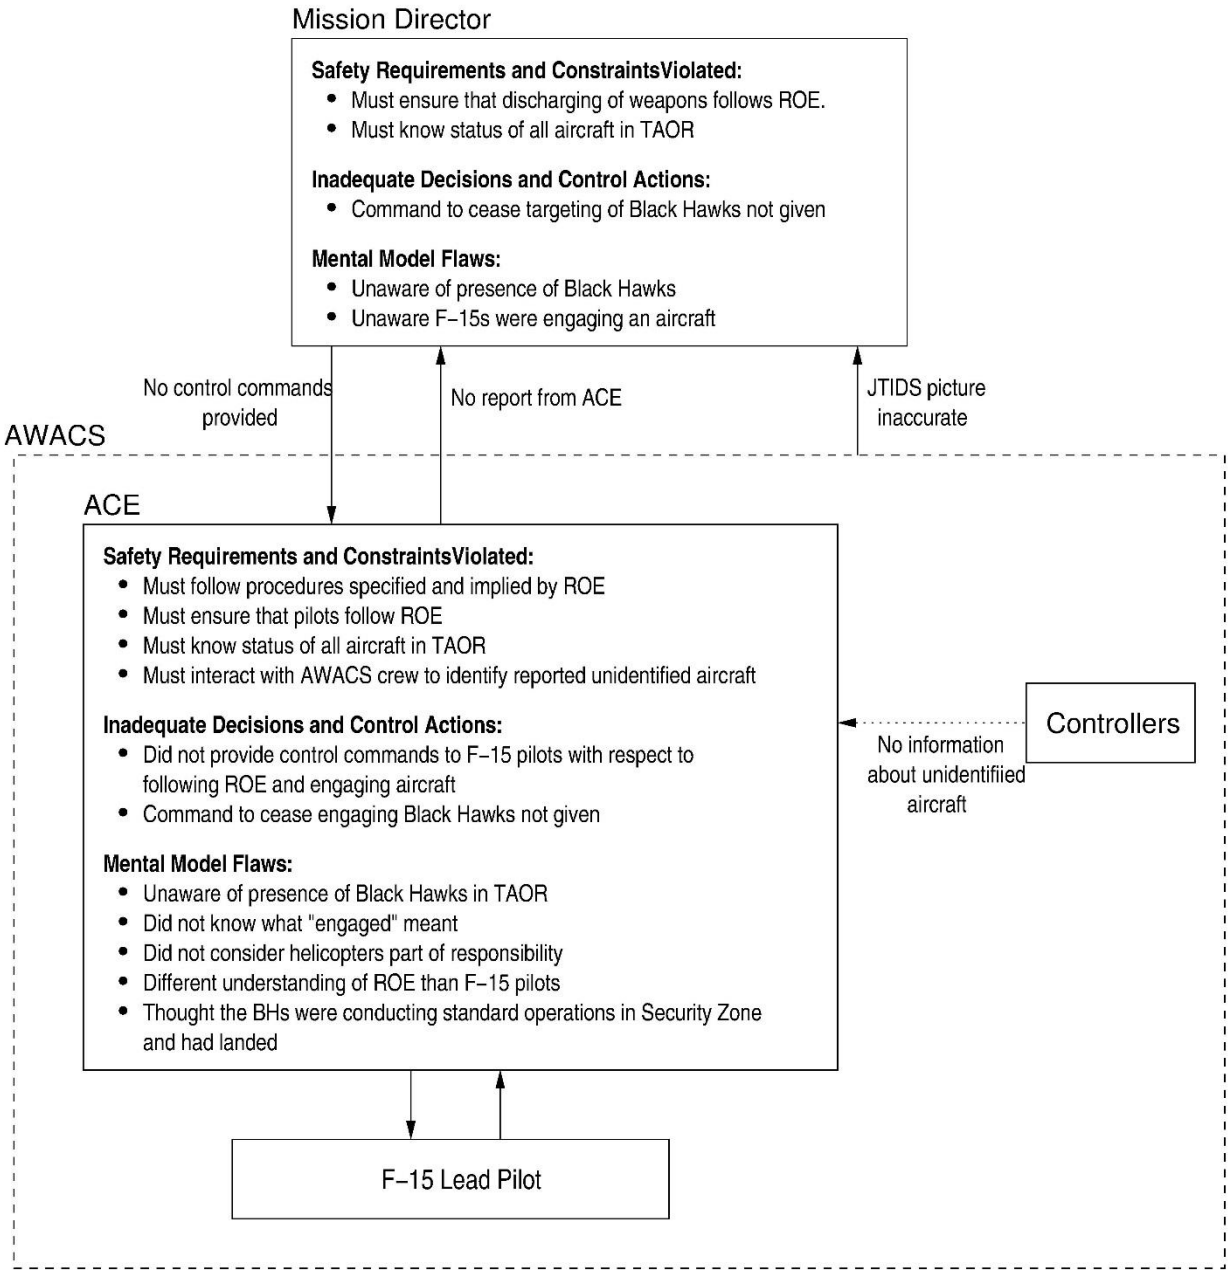
Mission Director
Safety Requirements and ConstraintsViolated:
Must ensure that discharging of weapons follows ROE.
Must know status of all aircraft in TAOR
Inadequate Decisions and Control Actions:
Command to cease targeting of Black Hawks not given
Mental Model Flaws:
Unaware of presence of Black Hawks
Unaware F–15s were engaging an aircraft
No control commands provided
No report from ACE
JTIDS picture inaccurate
AWACS
ACE
Safety Requirements and ConstraintsViolated:
Must follow procedures specified and implied by ROE
Must ensure that pilots follow ROE
Must know status of all aircraft in TAOR
Must interact with AWACS crew to identify reported unidentified aircraft
Inadequate Decisions and Control Actions:
Did not provide control commands to F–15 pilots with respect to following ROE and engaging aircraft
Command to cease engaging Black Hawks not given
Mental Model Flaws:
Unaware of presence of Black Hawks in TAOR
Did not know what "engaged" meant
Did not consider helicopters part of responsibility
Different understanding of ROE than F–15 pilots
Thought the BHs were conducting standard operations in Security Zone and had landed
Controllers
No information about unidentifiied aircraft
F–15 Lead Pilot

AWACS Mission Crew

Safety Requirements and ConstraintsViolated:

- Must identify and track all aircraft in TAOR
- Friendly aircraft must not be misidentified as hostile
- Must accurately inform fighters about status of all aircraft when queried.
- Must alert fighters of any aircraft not appearing on flowsheet.
- Must not fail to warn fighters about any friendly aircraft they are targeting
- Must provide ground with accurate picture of airspace and its occupants (through JTIDS).

Dysfunctional Interactions:

- Control of aircraft not handed off from enroute to TAOR controller
- Interactions between ASO and senior WD with respect to tracking the flight of the helicopters on the radarscope.

Inadequate Decisions and Control Actions:

- Enroute controller did not tell BH pilots to change to TAOR frequency.
- Enroute controller did not hand off control of BHs to TAOR controller
- Enroute controller did not monitor course of BHs while in TAOR.
- Enroute controller did not use Delta point system to determine BH flight plan
- TAOR controller did not monitor course of helicopters in TAOR
- Nobody alerted F–15 pilots before they fired that the helicopters they were targeting were friendly.
- Nobody warned pilots that friendly aircraft were in the area.
- Did not try to stop the engagement
- Nobody told BH pilots that squawking wrong IFF code.
- MCC did not relay information that was not on ATO about helicopters during morning briefing.
- Shadow crew was not monitoring activities.

Coordination Flaws:

- Confusion over who was tracking helicopters
- Confusion over responsibilities of surveillance and weapon directors
- No one assigned responsibility for monitoring helicopter traffic in NFZ
- Confusion over who had authority to initiate engagement

Context:

- Min Comm
- Poor morale. inadequate training, over worked
- Low activity at time of accident
- Terminal failure led to changed seating arrangement
- Airspace violations were rare

Mental Model Flaws:

- Did not think helicopters were an integral part of OPC air operations.
- Inaccurate models of airspace occupants and where they were.
- Thought helicopters only going to Zakhu

Systemic factors can often be described effectively with figures. The communication losses in the Überlingen accident was shown using two figures (Figure 16 and 17): one showed the designed communication links and the other showed the links actually operable at the time of the accident. Missing and flawed feedback was shown in Figure 18 for the ComAir Lexington crash. Figures can also be an effective way to show important changes over time in the control structure. Figure 20 shows the original designed control structure involved in an E. coli contamination of a water supply in Walkerton Canada. The details are not important here. Figure 21 shows the control structure as it existed at the time of the accident. The greyed-out structures in Figure 21 show the parts of the control structure that were eliminated over time with the blue structures depicting additions.

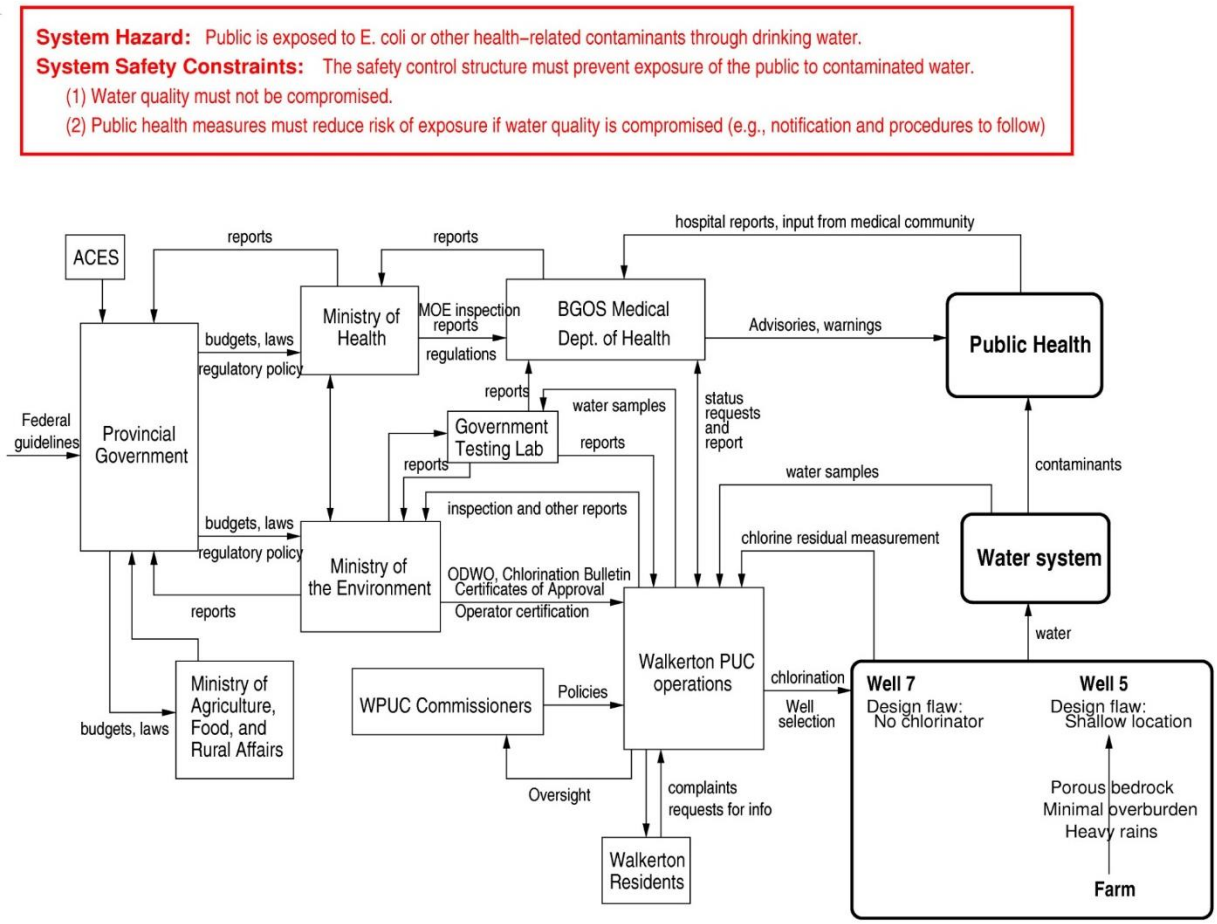

Figure 20: The original, designed control structure to control water quality in Ontario, Canada.

In this water contamination accident, a conservative government had been elected and decided to privatize the government water testing lab. This change would have been OK if the private labs had been required to inform the government about the results of their water testing. However, there was also an effort to reduce what was considered to be red tape and excessive reporting by private companies, so the private testing labs only reported their results directly to the operators overseeing water supply operations in Walkerton, who did not understand that water could be contaminated with E. coli. Figure 21 shows the resulting loss of all feedback about the behavior of the Walkerton operations by the Ministry of the Environment. By shading out the parts of the control structure that had disappeared, the problem is clearly depicted. While the Medical Dept. of Health might have been able to detect the problems, they assumed that the Ministry of the Environment was providing oversight and did not intervene. The blue parts of the control structure show the additions over time. Like any accident, the explanation is quite complicated. For one thing, although operator certification was instituted (which would have required additional education for the Walkerton operators), they were grandfathered out of new education and certification requirements.

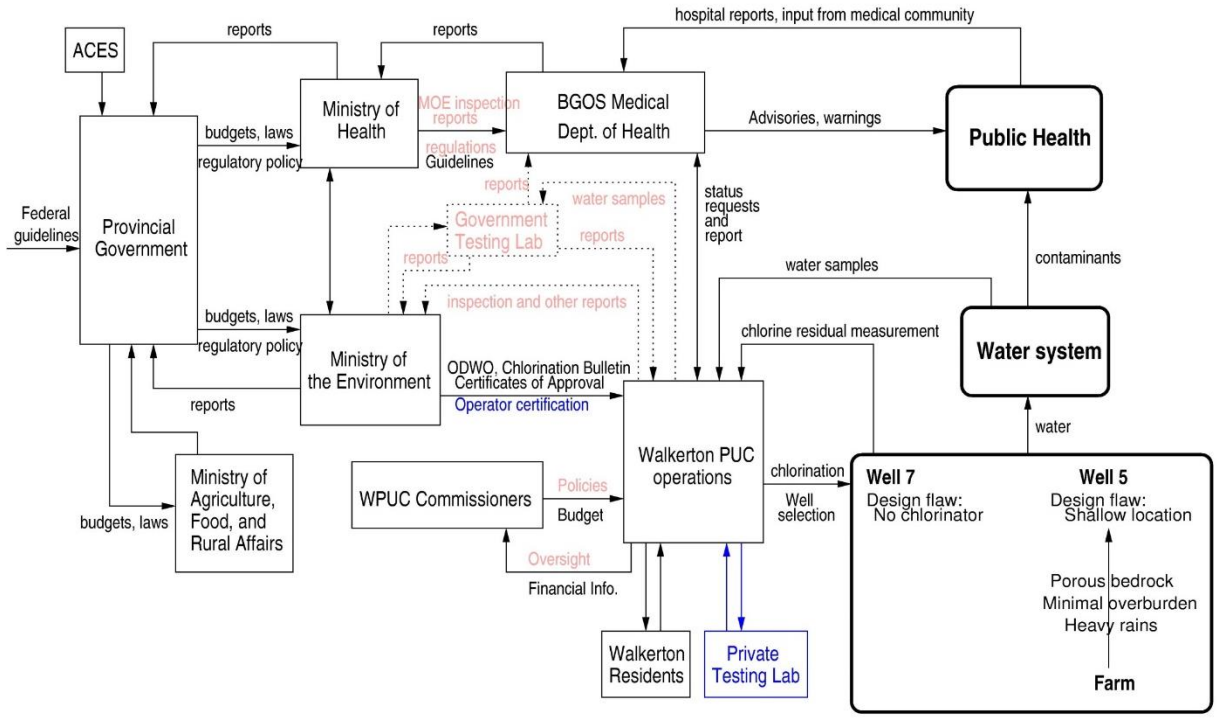

Figure 21: The control structure that existed at the time of the water contamination events.

Chapter 5: Using CAST for Workplace and Social Accidents

The examples so far in this handbook have involved traditional technical (engineered) systems plus the social systems that support them. CAST, however, can be used in less traditional applications. In this chapter, an example is provided of its use in workplace safety and also in situations where the hazards primarily involve social systems.

Workplace Safety

As the example in this section demonstrates, there is no difference in the way that CAST is used for workplace safety compared to other types of systems. The example was a real accident, but unimportant details have been omitted for confidentiality reasons. Unlike the other examples in this handbook, there was no public investigation of the events.

The accident occurred in a manufacturing plant where large products were being assembled. A part was not ready at the time it should have been and had to be installed later in the assembly process. To accomplish the goal, part of the structure of the product needed to be disassembled so the new part could be inserted. Scaffolding was needed to perform the disassembly and the scaffolding was set up by the first shift. In the second shift, a four-man team performed a pre-task analysis and devised a plan to remove part of the structure to get to the space where the new part was to be inserted.

When the plan was executed, it was quickly discovered that the product could not be disassembled because the scaffolding blocked the removal of the outside structure. The four-man team decided to remove some of the center planks of the scaffolding to make the necessary room.

The team disassembled the structure and started moving it toward the scissor lift that was positioned at the opposite end of the scaffolding. While moving the structural part, a mechanic fell through the hole in the scaffolding, sustaining a serious injury involving multiple broken ribs.

A standard workplace injury analysis was performed by the company with the following conclusions taken directly from the report written by their accident investigation team:

Direct Cause: Employee fell from a height.

Contributing Causes:

- Lack of experience performing this work (it was out of the usual sequence of activities).
- Workers had no knowledge that there was a shop aid used for this type of job.
- The team did not perceive an out of ordinary risk in performing this job and did not ask for help.

Root Cause: Floor boards removed from scaffolding

Near-Term Corrective Actions (Recommendations):

- Create an alert to communicate that floorboards cannot be removed from scaffolding i.e., tell workers that it is very important not to remove floorboards from scaffolding.
- Ensure all floorboards in the factory are secured. [In the CAST analysis it will be seen that there are safety reasons for <u>not</u> securing the floorboards.]

Longer-Term Corrective Actions:

- Add tool information to job instructions for out-of-sequence work.
- Reset level of acceptable risk using this and other examples [*What does this mean?*]
- During daily crew kickoff meetings, discuss potential hazards and ensure safe work practices for the assigned tasks of the day. Move away from talking about generic safety

topics. [*It seems that only with hindsight bias would someone think to tell the crew not to remove floorboards from scaffolding. If everything that could possibly be done wrong was covered in the briefing, such meetings would take all day*.]

During the investigation, the analysis was at the level of asking and answering questions like

Question: "What did the employee do wrong? Answer: "Stepped in a hole."

Question: "What were the possible causes?" Answer: "Lack of situation awareness. Made a mistake."

The causes were derived in this instance by using something called a Reality Chart, which appears to be simply a fault tree with a misleading name. There is no such thing as objective "reality" but only reality as seen by each individual—which will usually be different for each viewer of or participant in an activity. That "reality" may be biased by the assumptions and prejudices of those doing the analysis. It may also be incomplete. Labeling something as "reality" is simply a misleading sales tool.

In this case, "lack of situation awareness" is a technical term that means "the person did not know the hole was there when he stepped in it." That is pretty obvious and not helpful. Why would anyone purposely step in a hole? It also is probably wrong. The worker did know about the hole. So, we still do not know why he stepped in it. One guess is that the large structural part being moved involved four people holding different sections. The worker that fell may have been forced into stepping into the hole by the movement of the part by the other workers. Note that even knowing more about why the worker stepped in the hole, in this case, does not provide much guidance in preventing this accident. The second reason given, i.e., "made a mistake," is even less useful.

One thing to note is that, as is usually the case, the accident report focused on the victim and the lowest level workers directly involved with the job being performed.

Lest the reader is unfamiliar with workplace injury causal analysis and thinks this example is a particularly bad one, be assured that it is not atypical. The company involved is very large, produces highly complex systems, and the process used to analyze this accident followed their standard written procedures and the process used in many, if not most, companies.

The corrective actions (recommendations) do not seem unreasonable until one compares them with those created by CAST. Note that allowing floorboards to be removed is a requirement to accomplish other safety and manufacturing goals so the consequences may conflict with other goals and therefore is likely to be ignored, particularly over time. Adding tool requirements for specific jobs in the factory certainly seems like a good idea (and surprising that it is not already done) but the other long-term corrective actions seem less useful or realistic. And, again, everything comes down to the workers following written directions, some of which may be unsafe.

CAST Analysis of Scaffolding Accident:

Because so little information was contained in the official accident report, the CAST analysis will necessarily be filled mostly with questions that needed to be answered at the time. But here is an outline of the most basic system components and CAST-generated questions:

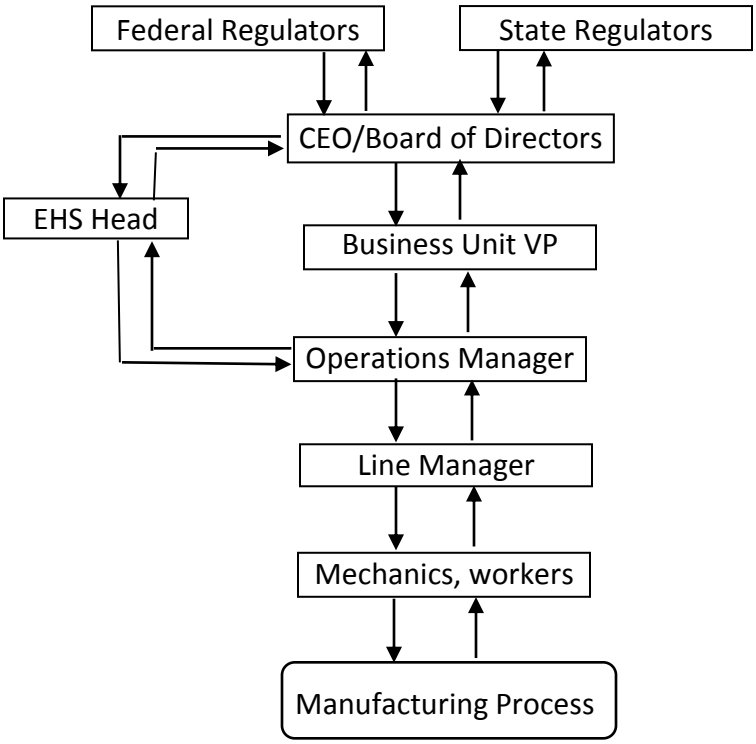

Physical:

Failures: none

Unsafe interactions: Worker fell through a hole in scaffolding

Workers on First Shift:

Responsibilities:

- Follow requests of their supervisor
- Follow documented standard procedures

Contribution to accident:

- Set up wrong type of scaffolding configuration

Mental Model Flaws:

- Did not know what job the scaffolding was for or did not understand the requirements of that job.

Context:

- Were they told to set up the scaffolding the way that they did?
- It was an out of sequence job that had not been completed at the time it should have been
- There was no documentation of the tools needed for the job or, alternatively, they may not have been told what the job was. [*What did they know about the job that needed the scaffolding and its requirements?]*

First Shift Supervisor:

Responsibilities:

- Provide job instructions to workers and ensure they have done the right job

Contribution to the accident:

- Either told workers to set up wrong type of scaffolding or did not check that right type was set up

Mental model flaws:

- *Did he know what job the scaffolding was needed for?*
- *Did he know there were special requirements for the scaffolding? If not, why not?*

Context:

- Out of sequence job that had not been completed at the time it should have been
- No documentation of tools needed for the job or they were not told what the job was.

Second Shift Workers:

Responsibilities:

- Perform their jobs as taught
- Use standard practices and tools called out in job instructions

Contribution to accident:

- Used unsafe scaffolding (removed the floor boards)
- *Did they check the jury rigging with their supervisor before they did it?*

Mental Model Flaws:

- Did not know there was scaffolding that should have been used for this job
- Did not know that they should not remove scaffolding to jury rig the tools they had been provided for the job.

Context:

- Unsafe scaffolding was set up for them to use that did not allow them to do the job they were told to do
- Correct tools were not documented
- *What experience did they have in doing this job? Was it the first time they had done it*? *If it was the first time, why were they unsupervised*?
- The reason for allowing scaffolding (floor boards) to be removed is that using wire or other materials to prevent scaffolding from being disassembled avoids the potential for FOD (foreign object debris), which is a serious hazard for these products.
- The culture in this company, which is documented in company policy, is that the highest priority is for workers to find a way to accomplish their goals without losing time in waiting for direction in how to do so from others. They are encouraged to find a way to get around problems without asking for instructions or help from others.

Supervisor on Second Shift:

Responsibilities:

- Ensure workers are using the proper and safe tools for job
- Ensure workers are aware of hazards, especially when doing out of sequence work

Contribution to accident:

- Allowed the use of unsafe scaffolding
- Did not properly supervise this out of sequence work

Mental model flaws:

- *Did he know the potential hazards?*
- *Did he know there was scaffolding that was designed for this type job? If not, why not?*

- *Did he know the workers had removed the floor boards?*
- *Did he check that proper scaffolding had been set up by prior shift?*

Context:

- Proper tools were not documented
- Out of sequence work is common in the factory

It is not possible to provide, after the fact, a detailed and useful CAST analysis of the upper levels of control in this accident because of the lack of information included in the accident report. There are, however, some intriguing unexplained statements in the accident report. For example, "The meeting with [company providing the out-of-sequence part] employees did not occur as planned." This statement was used in the explanation for why the wrong scaffolding was used. *Was there a meeting planned about how to fix the factory omission? Why did it not occur? Why did the work occur even though the meeting to plan how to do it was delayed?*

There are lots of questions that need to be answered to complete a CAST analysis and understand the cause of this accident, including:

- *Why were the workers doing a job for which they had no experience and without oversight from someone who did?*
- *Why did they not know about proper job aids, beyond the lack of documentation? What type of training is provided with respect to such jobs?*
- *Who has the responsibility to ensure that the right equipment is available and used? The immediate supervisor? Higher-level controllers?*
- *Why did the workers not ask for advice when the scaffolding design prevented them from doing their job? The answer here is probably related to the company culture.*
- *Who provides oversight for out-of-sequence work?*
- *Who evaluates the hazards of out-of-sequence work?*
- *Who is responsible for documenting what tools are needed for specific jobs and ensuring that workers know what these are? Why is it not standard practice to call out the tools to be used?*
- *How are out-of-sequence job steps and tools explained to employees?*
- *It appears that workers are used to jury-rigging solutions with no inputs from others. Is that true? How often does that occur? Why has nobody stopped this practice in the past? Is it expected by management?*
- *Why was the blame for not understanding the risks involved placed on the lowest-level workers and not their supervisors?*
- *Why was incorrect scaffolding constructed in the first place?*
- *Why was the part not inserted when it was originally supposed to be?*
- *There was a meeting scheduled about how to accomplish the work that had been canceled. Why was it never held? Why didn't the work wait until it could be held?*

- *Did higher-level management know about these practices? If so, why was nothing done about it before an injury occurred? If not, why not? Were there no performance audits or feedback to oversee how work was being accomplished in the factory?*
- *Out-of-sequence work was common and low-level supervisors worried about injuries resulting from it but they were provided with no assistance in dealing with it. Why not?*

Answering these questions and additional ones that arose in the process is likely to generate more recommendations than "Tell workers not to remove floorboards from scaffolding." Here are some preliminary recommendations, without completing a CAST analysis:

- Do not allow workers to do any job without someone supervising it who is familiar with the tools that should be used for it and without being given proper instructions.
- Control the safety of out of sequence work much better than is currently done.
- Reduce the incidence of out-of-sequence work. While it probably cannot be totally eliminated, the occurrence arises in problems in the supply chain that need to be tackled.
- Plan contingency actions for out of sequence work and make sure proper tools are available. At the least, someone should be supervising this type of work.
- Document the tools, including scaffolding, that should be used for every type of job, not just this one.
- Provide better feedback and performance auditing of how hazardous activities are being performed and how the hazards are controlled.
- Perform better workplace hazard analysis and control hazards that are identified.

One of the things I learned from studying workplace training at this company is that workers are expected to identify the hazards of their jobs themselves and to take steps to mitigate them. Surprisingly, I have found this same expectation at other companies. One even asked workers to do a human factors analysis of their own jobs. These are people who mostly do not have a college education and are not qualified to accomplish these tasks. Where are the safety professionals who should be doing these things? Making the workers totally responsible for their own safety, of course, does provide a convenient scapegoat when accidents occur.

CAST can easily be applied to workplace accidents without significant changes to the process.

Using CAST for Analyzing Social Losses

Most of the CAST analyses we have done ourselves have involved complex and dangerous physical systems and processes. CAST, however, can also be used to analyze social processes and other types of losses. This section provides some examples we have done. Links to the complete CAST analyses as well as many others are provided Appendix A. Two examples are described here: pharmaceutical safety in the U.S. and the causes of the 2008 financial system meltdown.

Pharmaceutical Safety Example

Vioxx (Rofecoxib) was a prescription pain management drug primarily used by patients suffering from osteoarthritis. It was approved by the Food and Drug Administration (FDA) in May 1999. While on the market, it was one of the major sources of revenue for Merck, estimated to represent 11% of Merck's

sales in 2003. In September 2004, after the U.S. drug safety control structure had allowed this dangerous drug to be left on the market for over five years, Merck voluntarily withdrew the drug because of safety concerns. According to an epidemiological study by an FDA scientist, Vioxx was associated with more than 27,000 heart attacks or deaths: "Vioxx may be the single greatest drug safety catastrophe in the history of this country or the history of the world."[23]

The losses in this example are identified as:

Losses ("Accidents"):

1. Patients get a drug treatment that negatively impacts their health, in this case fatal cardiovascular events such as heart attacks and strokes
2. Patients do not get the treatment they need.

Emphasis in this example is on the first loss, i.e., patients get a drug treatment that negatively impacts their health, because Vioxx, in the end, did not appear to have been more effective than normal over-the-counter pain medication so they could have gotten the treatment they needed.

Hazards: (common to all pharmaceutical products)

H1: The public is exposed to an unsafe drug

1. The drugs are released with a label that does not correctly specify the conditions for safe use of the drug
2. Approved drugs are found to be unsafe and appropriate responses are not taken (warnings, withdrawals from the market, etc.)
3. Patients are subjected to unacceptable risk during clinical trials

H2: Drugs are taken unsafely

1. The wrong drug for the indication is prescribed
2. The pharmacist provides incorrect medication
3. The drugs are taken in an unsafe combination
4. The drugs are not taken according to directions (dosage and timing)

H3: Patients do not get an effective treatment they require

1. Safe and effective drugs are not developed or are not approved for use
2. Safe and effective drugs are not affordable for those who need them
3. Unnecessary delays are introduced into development and marketing
4. Physicians do not prescribe needed drugs or patients have no access to those who could provide the drugs to them
5. Patients stop taking a prescribed drug due to perceived ineffectiveness or intolerable side effects.

The reader can see that some of the hazards conflict, making decision making about pharmaceuticals quite complicated and difficult.

These hazards can be translated into safety requirements and constraints:

[23] Testimony of David J. Graham, MD, MPH, November 18, 2004, U.S. Senate, 2004.

Safety Requirements and Constraints for Pharmaceutical Products:

1. Pharmaceutical products are developed to enhance long-term health
    a. Continuous appropriate incentives exist to develop and market needed drugs
    b. New scientific knowledge and technology is developed to create new drugs
    c. New drugs are developed and manufactured only when the scientific and technical knowledge is available

2. Drugs on the market are adequately safe and effective
    a. Drugs are subjected to effective and timely safety testing
    b. New drugs are approved by the FDA based upon a validated and reproducible decision-making process
    c. Drug approval is not unnecessarily delayed
    d. The labels attached to drugs provide correct information about safety and efficacy
    e. Drugs are manufactured according to Good Manufacturing Practices
    f. Marketed drugs are monitored for known and unknown adverse events, side effects, and potential negative interactions
    g. Long term studies are conducted, even after the drug as been approved, to validate the FDA's approval decision (e.g., Phase IV studies) both on the long term and for subpopulations
    h. New information about potential safety risks is reviewed by an independent advisory board
    i. Marketed drugs found to be unsafe after they are approved are removed, recalled, restricted, or appropriate risk/benefit information is provided

3. Patients get and use the drugs they need for good health
    a. Drugs are obtainable by patients
    b. Accurate information is available to support decision-making about risks and benefits
    c. Patients get the best intervention reasonable for their health needs
    d. Patients get drugs with the required dosage and purity

4. Patients take the drugs in a safe and effective manner
    a. Patients get correct instructions about dosage and follow them
    b. Patients do not take unsafe combinations of drugs
    c. Patients are properly followed by a physician while they are being treated
    d. Patients are not subjected to unacceptable risk during clinical trials

The safety control structure we used has five main components (Figure 22): the FDA, Pharmaceutical companies, patients, academically-affiliated researchers, and the healthcare providers as well as a variety of others who play a role in drug safety (Congress, Journal editors, patient groups, etc.). The responsibilities of the controllers are described in the individual component CAST analyses. Only a few

examples from the CAST analysis are presented here. The complete analysis by Matthieu Couturier can be found in Appendix A.

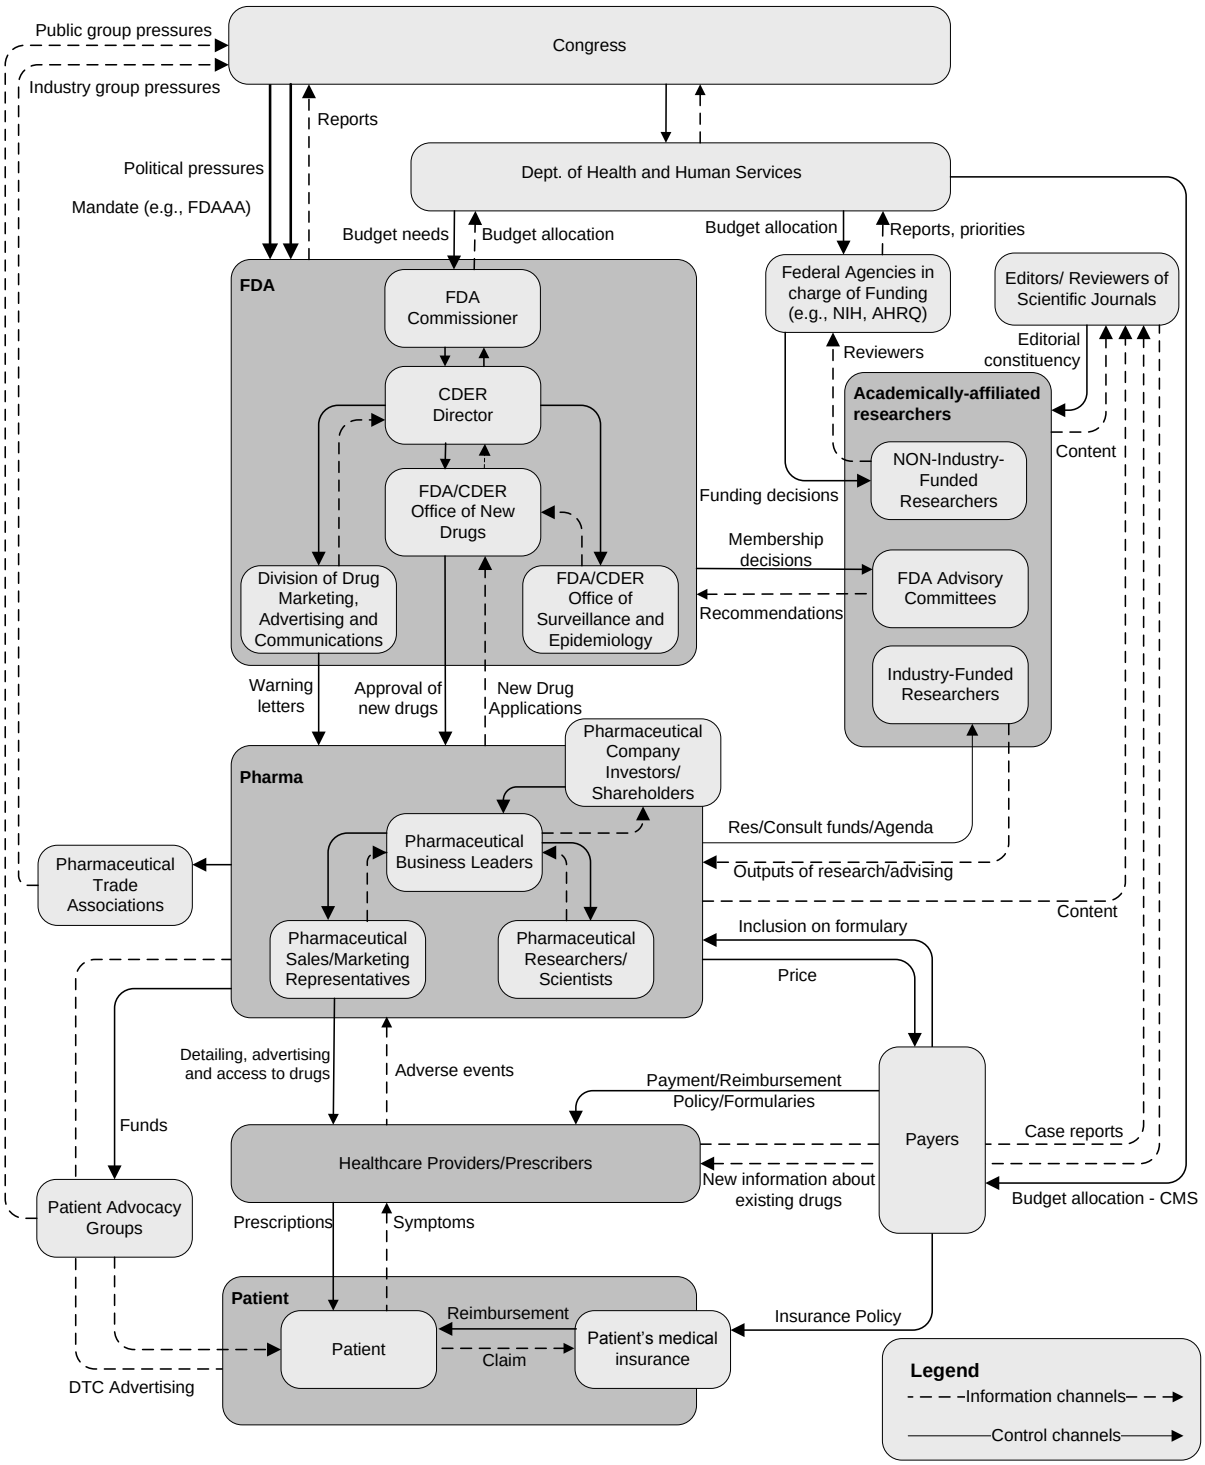

Figure 22: The pharmaceutical safety control structure in the U.S.

Patients

Responsibilities:

1. Accept limited responsibility for their own health and treatment (limited by what is practical).
2. Follow their physician's instructions and take drugs as prescribed.
3. Accede to doctor's superior knowledge when appropriate.
4. Go through a doctor to get a prescription for drugs like Vioxx.

Contribution to Losses:

1. Some patients pressured their doctor into prescribing Vioxx even if it was not necessarily the most appropriate treatment for their specific needs.

Mental Models:

1. Patients believed that the drug was safer than it really was.
2. Patients believe that newer, more expensive drugs are better than older, alternative treatments.

Context in Which Decisions Were Made:

1. Patients had limited information about the safety and effectiveness of Vioxx. Most of the information that they had came from Direct-To-Consumer (DTC) advertising, which provided a rosy picture of the efficacy of the drug along with glamorous and respected spokespersons (e.g., the famous ice skater Dorothy Hamill).
2. Patients have limited medical knowledge about both their disease and the medication they are taking and have to rely on their doctor's opinion and advice.
3. Vioxx was approved by the FDA which they assumed provided a "guarantee" of safety. It was apparently endorsed by the medical community in medical journals and tacitly endorsed by insurance companies (which are willing to reimburse the patients for the new treatments).
4. A large number of patients taking Vioxx should not have been taking the drug because they had very low risk of gastrointestinal problems and therefore should have been prescribed non-steroidal anti-inflammatory drugs (NSAIDs).

Healthcare Providers/Prescribers

Responsibilities:

1. Make treatment decisions based on the best interests of their patients.
2. Weigh the risks of treatment and non-treatment.
3. Prescribe drugs according to the limitations on the label.
4. Maintain an up-to-date mental model of the risk/benefit profile of the drugs they are prescribing.
5. Monitor symptoms of their patients under treatment for adverse events and negative interactions.
6. Report adverse events potentially linked to the use of the drugs being prescribed.

Contribution to Losses:

1. Doctors prescribed Vioxx, both on and off label, for patients for whom it was not indicated.

Mental Models:

1. Believed that new drugs are better than existing treatments.
2. Believed that information from pharmaceutical companies is accurate.
3. Did not understand the risk/benefit tradeoffs of Vioxx. In particular, they did not know about potential cardiovascular risks associated with the long-term use of the drug.
4. Believed that patients might go to another practitioner if Vioxx was not prescribed.

Context in Which Decisions Were Made:

1. Doctors mostly learn about new products from the drug companies themselves.
2. Doctors are notoriously busy and their time is limited for reading medical journals and keeping up with the latest treatments or with thoroughly discussing the benefits and risks of a treatment.
3. Doctors have limited access to unbiased information.
4. Studies of new drugs are typically done against placebos. Doctors are left without the information needed to decide which approved drug is a more appropriate treatment for a specific patient.
5. Doctors are part of the service industry and do not want to alienate their patients by not prescribing the drugs they request.
6. Vioxx label (from FDA) did not mention cardiovascular risks.

Merck

Responsibilities:

1. Ensure that patients are protected from avoidable risks:
   a. Provide safe and effective drugs.
   b. Test drugs for effectiveness.
   c. Properly label the drugs.
   d. Protect patients during clinical trials by properly monitoring the trial.
   e. Do not promote unsafe use of the drugs.
   f. Remove a drug from the market if it is no longer considered safe.
   g. Manufacture the drugs according to Good Manufacturing Practices.
2. Monitor drugs for safety:
   a. Run long-term post-approval studies as required by the FDA.
   b. Run new trials to test for potential safety hazards.
   c. Provide, maintain, and incentivize adverse-event reporting channels.
3. Give accurate and up-to-date information to doctors and the FDA about drug safety:
   a. Educate doctors.
   b. Provide all available information about the safety of the drug to the FDA.
   c. Inform the FDA of potential new safety issues in a timely manner.
4. Conduct or sponsor research that can be useful for the development of new drugs and treatments.

Contribution to Losses:

1. Merck did not run studies that might have found negative cardiovascular (CV) results. Company executives rejected doing a study of Vioxx's CV risks.
2. Merck's studies, and the results the firm published, inadequately represented the risk/benefit profile of the drug:
   a. The studies were motivated by marketing goals.
   b. If the results were published, they were typically released very late or only partially released.
   c. The results were biased to appear better than they were.
   d. Some of the studies that were run did not have an active Data and Safety Monitoring Board (DSMB) to monitor the clinical trials and protect the patients. The safety of the patients was solely in the hands of the Merck investigators.
3. Merck published and disseminated misleading information about the safety profile of Vioxx:
   a. Merck aggressively promoted drug usage with a task force trained to avoid CV questions.
   b. Merck used promotional activities and materials that were false, lacking in fair balance or otherwise misleading. The company continued to minimize unfavorable findings up to a month before withdrawing Vioxx.
   c. Merck published or promoted publication using guest authorship or ghostwriting; Merck employees' involvement in the writing process was often not mentioned and the financial ties of the main authors were not always disclosed.
   d. Merck created journals made to look like independent peer-reviewed journals. These journals were actually marketing compilations of articles promoting Vioxx and another drug made by Merck.

Mental Models

1. Merck originally believed that Vioxx did not cause any CV events.
2. Believed Vioxx had the potential to become a blockbuster drug and significant source of revenue if only Merck could effectively limit the publication of negative results and convince doctors to prescribe the drug despite potential CV risks.
3. Believed the company could protect its reputation by hiding negative results.
4. Believed it could convince doctors to prescribe the drug despite the potential CV risks.

Context:

1. Merck has a fiduciary duty to shareholders to provide a return on their investment and stakeholders demand a high return. Furthermore, drug company executives are partly paid in stock options.
2. Satisfactory financial results depended on Vioxx being a blockbuster.
   a. Vioxx was a major source of Merck's revenue and was extremely profitable.
   b. Vioxx faced fierce competition from Celebrex, which was approved by the FDA 5 months before it approved Vioxx. Merck had to aggressively promote the drug to be competitive with Celebrex.
   c. Merck could not allow negative study results to impact sales.
   d. Pharmaceutical companies often depend on developing blockbuster drugs. The success of Vioxx was considered crucial for the financial future of the company.
3. Comparative studies suggested that Vioxx had a higher number of CV events than Naproxen. Merck assumed that difference came not from Vioxx having any negative side effects but rather because Naproxen protected patient's hearts.
4. Drug companies have no incentives to do Phase IV (post approval) studies or to publish negative results from internal studies.
5. Most clinical research on drugs is sponsored by companies that make them. Drug companies now have more control than in past on the way the research is carried out and reported.
6. Merck had a reputation to maintain. Withdrawing Vioxx from the market and acknowledging CV events would have hurt their reputation.
7. The drug pipeline was dwindling and older drugs were going off patent protection. Merck was about to lose five of its most profitable patents.
8. Drug companies have no incentive to do Phase IV safety testing, even if it is required by the FDA. Similarly, they have no incentives to publish negative internal studies.

FDA/CDER

Background Information

The section within the FDA responsible for human drugs is called the Center for Drug Evaluation and Research (CDER). Within CDER, the Office of New Drugs (OND) is in charge of approving new drugs, setting drug labels and, when required, recalling drugs. The Office of Surveillance and Epidemiology (OSE) focuses on identifying adverse events that were not detected during the approval of drugs and can recommend actions, such as label changes or recalls, to OND.

Responsibilities:

Committee Staff

1. Select competent advisory committee members and establish and enforce conflict of interest rules.
2. Provide researchers access to accurate and useful adverse events reports.

Office of New Drugs (OND)

1. Oversee all U.S. human trials and development programs for investigational medical products to ensure safety of participants in clinical trials. Provide oversight of IRBs (Institutional Review Boards) that perform these functions for the FDA.
2. Set the requirements and process for the approval of new drugs.
3. Critically examine a sponsor's claim that a drug is safe for intended use. Impartially evaluate new drugs for safety and efficacy and approve them for sale if deemed appropriate.
4. Upon approval set the label for the drug.
5. Do not unnecessarily delay drugs that may have a beneficial effect.
6. Require phase IV safety testing if there is a potential long-term safety risk.
7. Remove a drug from the market if new evidence shows that the risks outweigh the benefits.
8. Update the label information when new information about drug safety is discovered.

Division of Drug Marketing, Advertising, and Communications (DDMAC)

1. Monitor the marketing and promotion of drugs. Review advertisements for accuracy and balance.

Office of Surveillance and Epidemiology (OSE)

1. Conduct on-going reviews of product safety, efficacy, and quality. Perform statistical analysis on adverse event data received to determine whether there is a safety problem.
2. Re-assess risks based on new data learned after a drug is marketed and recommend ways to manage risk.
3. Serve as consultants to OND with regards to drug safety issues.
4. Recommend that a drug be removed from the market if new evidence shows significant risks.

Role Played:

Agency Wide

1. Allowed waivers of conflict of interest rules for advisory panel members.
2. Pressured an FDA employee to change conclusions and recommendations on a study of Vioxx and prevented publication of the results.
3. Was not able to provide quality adverse event reports for researchers to use.

OND

1. Gave expedited review and approval to Vioxx even though the drug did not meet the criteria for expedited review.
2. Did not check whether clinical trial safety requirements were being enforced (e.g., that protocol 078 had an active DSMB).
3. Approved Vioxx without requiring a Phase IV study even though the long term risks were unclear.
4. Did not update the Vioxx label in a timely fashion.
5. Delayed the recall of Vioxx. Did not act fast or effectively enough.

DDMAC

1. Original warning letter was not followed by subsequent warnings; false and misleading promotional material went un-reviewed.

OSE

1. Did not properly monitor the drug for long-term risks. Could not differentiate normal adverse events from the ones due to the drug within the population affected.
2. Did not insist that Merck launch a large-scale clinical trial when suspicions first arose.
3. Did not require a recall of the drug.

Context in Which Decisions Were Made:

Agency Wide

1. Lack of strong leadership at the head of the FDA, high turnover and unfilled positions.
2. Tensions between the Office of New Drugs (OND) and the Office of Surveillance and Epidemiology (OSE). OND was much larger than OSE and had more resources. When OSE proposes to recall a drug, it is perceived as a failure of OND.
3. PDUFA: The FDA is partly sponsored by the pharmaceutical companies it is supposed to monitor and regulate. This leads to (a) pressure to reduce approval time and (b) comparatively less staff in safety monitoring and marketing oversight than in approval.
4. Limited resources in personnel and budget.
5. Many of the experts who are asked to be on an advisory panel work with the pharmaceutical companies. It is difficult to find experts who do not have such ties.
6. Political and congressional pressures (e.g., anti-regulatory climate, special interests).

OND

1. For pre-market review, the FDA only has information provided by the company with no other independent research data.
2. Legislation inhibits full disclosure of information: Clinical information is kept secret and the FDA cannot share proprietary data for independent analysis.
3. The FDA is unable to keep track of ongoing clinical trials.
4. Legislation makes it difficult for the FDA to require label changes.
5. A very high certainty that the drug is dangerous is required before the drug is recalled.
6. OND is in charge of both approving and recalling drugs.
7. PDUFA fees represent more than 50% of ODN's budget; The FDA depends on PDUFA funding which affects the decision-making process.

OSE

1. No independent decision-making responsibility.
2. High turnover of OSE directors.
3. No control over selective publication (companies are not required to publish results from clinical trials or even the fact that clinical trials were being conducted).
4. No legal authority to require additional safety studies once a drug is approved. Post-marketing safety requirements are rarely enforced.
5. Adverse event reporting is limited: reporting voluntary, no certainty the event was actually due to the product, reports often not detailed enough for use. Researchers do not know how many people are taking the medication so cannot use the Adverse Event Reporting System' (AERS) data to calculate incidence in the population.
6. Very limited sources of information about adverse events.

The rest of the analysis is omitted here but a link to the thesis on this topic is in Appendix A.

2008 Financial Crisis Example

Melissa Spencer, in a 2012 MIT Master's thesis, demonstrated how CAST could be used to understand the collapse and rapid acquisition of the investment bank Bear Stearns in March 2008. Only a small part of the CAST analysis is included here. See Appendix A for a link to the complete thesis.

Loss Event: Bank Insolvency

Definition: Solvency is defined as satisfying the basic equation Assets = Liabilities + Equity. Insolvency occurs when this equation does not hold.

Hazards:

H1: Bank cannot pay debts (illiquid assets)
- Safety Constraint: Bank must be able to pay debts as they come due
- Safety Constraint: Bank must have adequate short-term assets and/or capital to cover short-term liabilities.

H2: Bank has greater liabilities than assets less equity
- Safety Constraint: Bank must always have a solvent balance sheet
- Safety Constraint: Bank must not acquire liabilities that exceed its assets and capital

H3: Market is unable to determine value of a financial instrument
- Safety Constraint: Market Participants must be able to determine value of market instruments
- Safety Constraint: All financial instruments traded in the market must have adequate information available to determine their inherent relative value.

H4: Market is unable to determine creditworthiness of institution
- Safety Constraint: Market Participants must be able to determine creditworthiness of bank
- Safety Constraint: There must be a clear and unbiased source of information about Bank's health available to the market

The primary components involved in the loss were:

- Asset Contributors
  - Traders at firm (buy and sell assets on the market)
  - Creditors to the firm (lend money or assets to firm)
  - Clients of the firm (provide fees and assets)
- Liability Controllers:
  - Management of Bear Stearns (responsible for protecting shareholders and managing risk
  - Traders at Bear Stearns (responsible for day to day balance sheet)
- Equity Contributors: Shareholders
- Federal Reserve Bank:
  - Overseas markets, sets interest rates, serves as lender of last resort to commercial banks
- Securities and Exchange Commission
  - Monitors balance sheets and health of investment banks
  - Monitors details of securities on the market
- U.S. Department of Treasury

Control Structure:

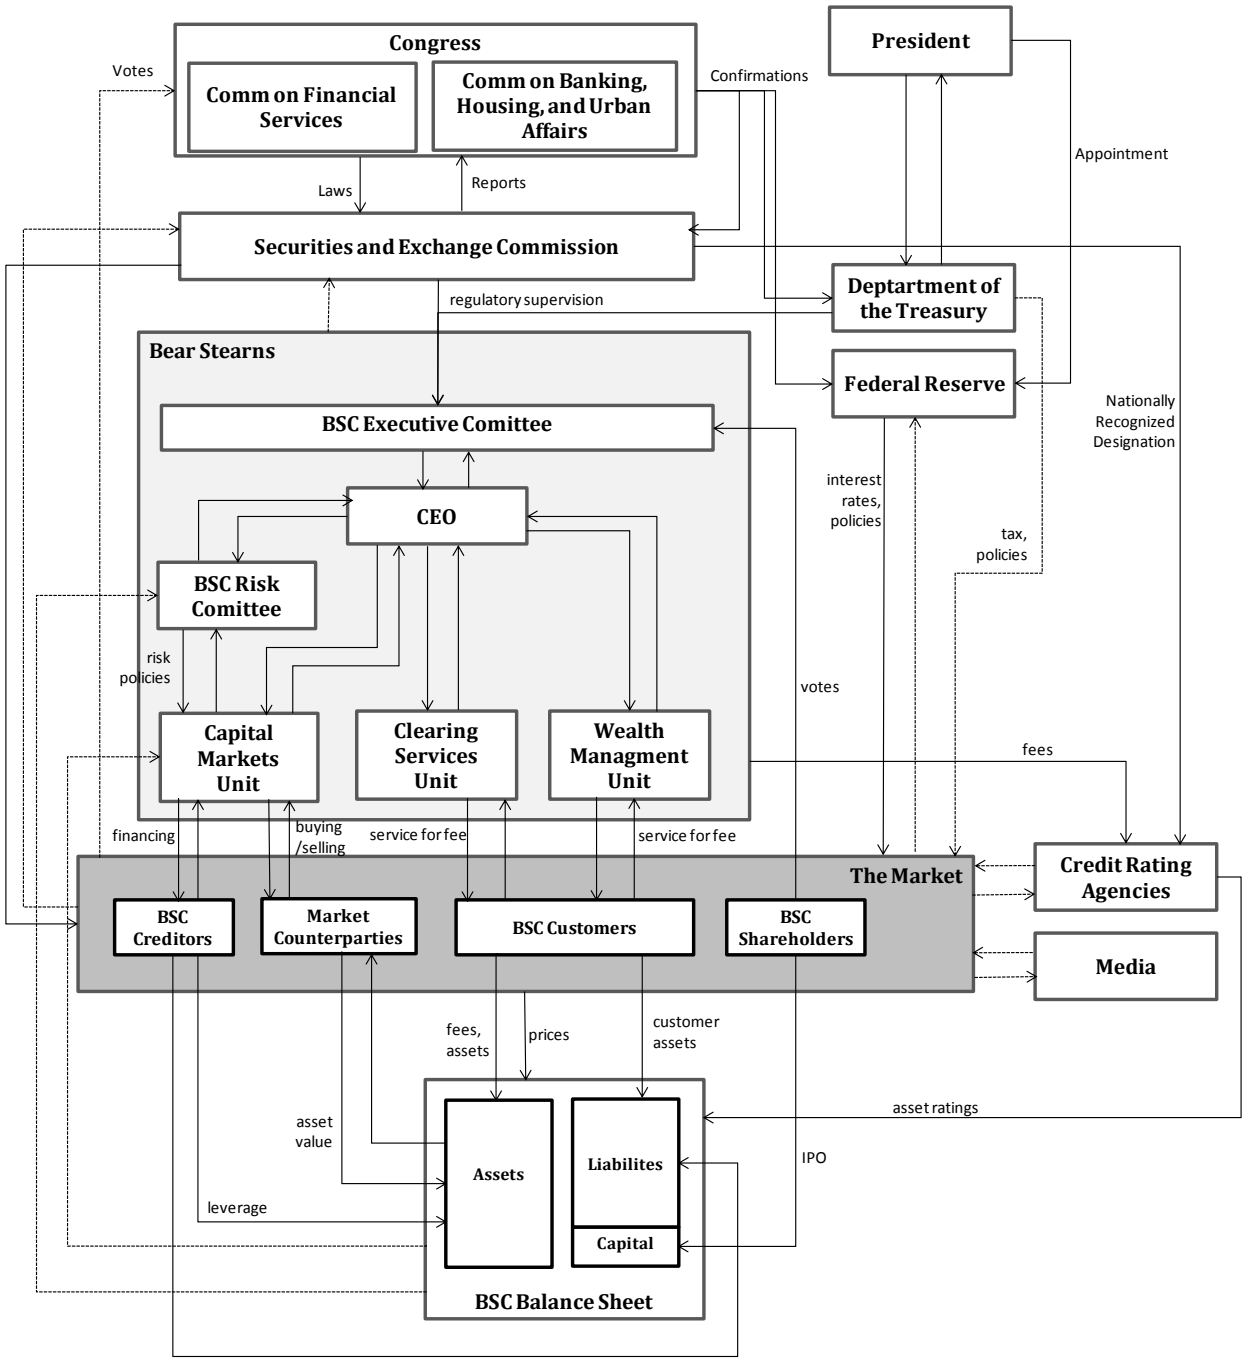

Only one component of the CAST analysis is provided as an example here:

Federal Regulators of Investment Banks (Federal Reserve):

Responsibilities:

- Ensure banks are free of fraud or deceitful practices.
- Ensure banks are following basic practices to protect customers.
- Ensure adequate mechanisms for public knowledge of investments, investors have access to adequate knowledge of risks, conduct audits of bank practices, monitor markets.
- Ensure national economy is safe, monitor interconnectedness of financial system, ensure all institutions and instruments are properly regulated, protect federal resources.
- Detect and eliminate market externalities that arise in rapidly changing markets.

Contribution to losses:

- Allowed regulatory distinction between commercial and investment banks to erode without a corresponding change to regulatory protections for both institutional types.
- Did not distinguish between health of system and health of bankers/shareholders within system
- Allowed interconnectedness between regulated and unregulated entities to build.
- Did not ensure there was an unbiased (not paid for by banks) credit rating mechanism.
- Did not update regulatory requirements for investment bank balance sheet quickly enough
- Promoted housing policies that led to unsafe mortgages.
- Did not require disclosure from independent credit rating agencies on operations.

Process Model Flaws:

- Believed financial institutions would “self-regulate” in order to stay competitive.
- Believed failure of unregulated and lightly-regulated entities would only impact investors directly exposed to them.
- Believed mortgage-based securities were created safely, reflecting accurate risk profiles to investors.
- Believed failure of a single institution or instrument would not impact the entire financial system.

Context:

- Complicated regulatory scheme (Fed, SEC, FDIC, US Treasure, etc.) with different jurisdictions
- Political pressure for a “laissez-faire” economic policy
- Commercial banks seen as “safe" (given FDIC and Fed regulations) while investment banks were for more savvy investors with greater risk appetite
- Growing political and social power of banks, who generally support easing regulation
- Vast differences in compensation between regulators and those they regulate
- Rapidly changing market (new technologies, new instruments, new connections) was very difficult for regulators to keep up with
- Banks constantly looking for (and finding) regulatory loopholes
- Political pressure to encourage home ownership in 1990s and 2000s.

There are only a few attempts so far to use CAST for non-physical systems, but it appears that it can be done.

Chapter 6: Introducing CAST into an Organization or Industry

CAST involves a paradigm change in the way people think about and identify accident causes. Introducing it into an organization may require greater effort than simply changing notations or making small changes in the old paradigm. A few ideas to assist in making this change are presented here.

Paradigm changes of this sort often require enough dissatisfaction and frustration with the status quo by important participants in the organization or industry that natural lethargy can be overcome. A major accident involving large costs may assist in this process. But waiting for a major loss is clearly not the ideal way to introduce changes in how accident investigation is done. Demonstrations of how much more effective the new approach is when applied to previous accidents may be very helpful in both demonstrating practicality and benefits.

This approach is particularly effective when applied to losses in which CAST uncovers causes not found originally. Jon Hickey, a graduate student who was in the U.S. Navy, noticed that the Coast Guard experienced an increased aviation mishap rate where seven Class A mishaps occurred in the 22-month period between 2008 and 2010. The mishap investigations had centered on the pilots using HFACS, an accident causal analysis technique based on Reason's Swiss Cheese Model. The investigations did not find common contributing or causal factors. Commodore Hickey applied CAST, focusing on the systemic factors instead of human errors. Using CAST, he was able to identify the common factors leading to the losses. The HFACS analysis had stopped at identifying symptoms and did not identify the common systemic causes that were important in all the accidents.[24]

Introduction of new concepts will invariably lead to resistance in conservative, government organizations that are reluctant to introduce new ideas. Attempts to change established accident investigation processes have commercial, training, and contractual cost implications. The key in successfully introducing changes lies in being able to demonstrate that a new process is time efficient, is cost effective, and, most important, provides results that can significantly improve future safety. One potentially effective approach is to compare the results of accident reports published by well-regarded groups, such as the NTSB, based on a linear event model with a reanalysis of the same accident using a systems approach like CAST. We have been doing these types of comparisons, but more are still needed.

A powerful set of constituents can be created when the benefits of an explanatory vs. accusatory approach that does not focus on blame can be demonstrated. Groups, such as pilots who get the brunt of blame for accidents, are discovering the power of this new approach and providing support. Important player groups, such as unions, user-groups, and professional organizations that can exert pressure for change, may also be influential.

As with any paradigm change or cultural change, there must be buy-in at the top. Without high-level leadership, it is difficult to implement changes suggested at lower levels. Leadership willing to grab the opportunity to make changes after a major loss will be particularly effective. It is sad, but unfortunately too often true, that significant change is not possible without an event-driven disruption of complacency. More optimistically, significant changes and improvements can be driven by strong leaders. Paul O'Neill at Alcoa is an impressive example where he was able to both greatly reduce accidents while, at the same time, increase profits by focusing on safety as a primary goal. Basically, someone in authority has to decide that reducing losses is more important than continuing to do what they have always done.

---

[24] Jon Hickey, A System Theoretic Analysis of U.S. Coast Guard Aviation—CG-6505 Mishaps, STAMP Workshop, March 26, 2013, http://psas.scripts.mit.edu/home/

A significant demonstration by smaller players in the industry may also be effective. CAST is being adopted by some smaller accident investigation authorities. Success shown by these efforts may lead others to follow.

There are additional practical questions that need to be answered. One is the cost of training people to think in a different way and use a new approach. Not everyone finds system thinking natural or easy. For a large organization that does its own internal investigation of incidents and accidents, a special group that is experienced in using CAST will get the best results. The investigation team should be independent of the management of the group in which the events occurred and report to a higher level of management. Having a trained group in the safety organization may be the most practical approach.

Appendix A: Links to Published Examples of CAST Analyses of Real Accidents

Most of the analyses included here are by the author or her graduate students, with a couple of exceptions. Many were produced some time ago and may not include some of the more recent advances in CAST development. They do, however, at least show how CAST can be used in a variety of industries. I have not included the hundreds of analyses that have been done by students in my MIT graduate classes as they were limited in time and information and were just learning the technique. I also have not included industrial examples undertaken by others because of the proprietary nature of such accident analyses and privacy concerns.

CAST Analysis of the Shell Moerdijk Accident by Nancy G. Leveson

Increasing Learning from Accidents: A Systems Approach Illustrated by the UPS Flight 1354 CFIT Accident by Shem Malmquist, Nancy Leveson, Gus Larard, Jim Perry, and Darren Straker, May 2019

The Underestimated Value of Safety in Achieving Organization Goals: CAST Analysis of the Macondo Accident. by Maria Fernanda Tafur Munoz, MIT Engineering and Management Master's Thesis, June 2017.

Systems-Theoretic Accident Model and Processes (STAMP) Applied to a U.S. Coast Guard Buoy Tender Integrated Control System. by Paul D. Stukus, MIT SDM Masters Thesis, June 2017

Learning from Accidents that are a consequence of complex systems. by John Thomas and Shem Malmquist, ISASI Conference

A Systems Approach to Analyzing and Preventing Hospital Adverse Events by Nancy Leveson, Aubrey Samost, Sidney Dekker, Stan Finkelstein, and Jai Raman. *Journal of Patient Safety* , in press, 2016

A STAMP Analysis of the LEX Comair 5191 Accident , by Paul S. Nelson, Master's Thesis, Lund University, Sweden, June 2008, supervised by Prof. Sidney Dekker.

Safety-Guided Design Analysis in Multi-Purposed Japanese Unmanned Transfer Vehicle. by Ryo Ujiie, System Design and Management Master's Thesis, September 2016.

Systems Theoretic Process Analysis Applied to an Offshore Supply Vessel Dynamic Positioning System. by Blake Ryan Abrecht, MIT M.S. in Engineering Systems Thesis, June 2016.

Systems Theoretic Accident Analysis of an Offshore Supply Vessel Collision. by John Michael Mackovjak, Master of Science in Technology and Policy, MIT, June 2016.

STAMP applied to Fukushima Daiichi nuclear disasteer and the safety of nuclear power plants in Japan. by Daisuke Uesako, MIT Master's Thesis, System Design and Management Program, June 2016.

System Theoretic Safety Analysis of the Sewol-Ho Ferry Accident in South Korea , by Yisug Kwon, MIT Master's Thesis, December 2015.

A Systems Approach to Patient Safety: Preventing and Predicting Medical Accidents Using Systems Theory, by Aubrey Samost, MIT Master's Thesis, June 2015

Comparison of SOAM and STAMP for ATM Incident Investigation by Richard Arnold, Master's Thesis, Lund University, Sweden, 2009, supervised by Prof. Sidney Dekker.

A CAST Analysis of a U.S. Coast Guard Aviation Mishap , by Jon Hickey, MIT Master's Thesis, May 2012, supervised by Dr. Qi van Eikema Hommes.

Application of CAST to Hospital Adverse Events , by Meaghan O'Neil, MIT Master's Thesis, May 2014

Application of CAST and STPA to Railroad Safety. , by Airong Dong, MIT Master's Thesis, May 2012

Engineering Financial Safety: A System-Theoretic Case Study from the Financial Crisis , by Melissa Spencer, MIT TPP (Technology and Policy Program) Master's Thesis, May 2012

A Systems Theoretic Application to Design for the Safety of Medical Diagnostic Devices , by Vincent Balgos, MIT SDM Master's Thesis, February 2012, supervised by Dr. Qi van Eikema Hommes

Application of a System Safety Framework in Hybrid Socio-Technical Environment of Eurasia. by Azamat Abdymomunov, MIT SDM Thesis, 2011. This thesis won the "Best SDM Master's Thesis" award at MIT.

A System Theoretic Safety Analysis of Friendly Fire Prevention in Ground Based Missile Systems, by Scott McCarthy, MIT SDM Master's Thesis, January 2013

Accident Analysis and Hazard Analysis for Human and Organizational Factors by Margaret Stringfellow, October 2010

A System Theoretic Analysis of the "7.23" Yong-Tai-Wen Railway Accident . This paper, by Dajiang Suo from the Computer Science and Technology Dept., Tsinghua University, Beijing, China, was presented at the 1st STAMP/STPA Workshop held at MIT on April 26-28, 2012

A Case Study of Vioxx using STAMP by Matthieu Couturier, MIT Technology and Policy Master's Thesis, June 2010.

There are also several full examples in Nancy Leveson, *Engineering a Safer World*, MIT Press, 2012:

- Friendly Fire Shootdown of a U.S. Blackhawk Helicopter by a U.S. F-15 over the Iraqi No-Fly Zone, pp. 103-167
- The explosion of a Titan IV booster launched from Cape Canaveral in 1999, pp. 469-493
- The E. Coli Bacterial Contamination of a Public Water Supply in Walkerton, Canada, pp. 497-515

Finally, there is an academic paper I wrote about accident analysis that may interest some people:

Applying Systems Thinking to Analyze and Learn from Events by Nancy Leveson, *Safety Science*, Vol. 49, No. 1, January 2010, pp. 55-64

Appendix B: Background Information and Summary CAST Analysis of the Shell Moerdijk Loss

Background:

On 3 June 2014, an explosion and fire occurred at the Shell Moerdijk plant in The Netherlands. Shell Moerdijk produces chemicals, such as ethylene and propylene, used to manufacture plastic products. Heat is first used to convert gasoil, naphtha, and LPG into a wide variety of chemicals. These chemicals are then used, among other things, as raw materials to produce other products at Shell Moerdijk, including those produced by the styrene monomer and propylene oxide (MSPO) plant involved in the accident.

Shell has two MSPO plants in Moerdijk: MSPO1 (commissioned in 1979) and MSPO2. The accident took place in the MSPO2 plant, which was designed in 1996 by the predecessor of what is now called Shell Projects and Technology,[25] the license-holder for the process. On the basis of a user agreement, Shell Moerdijk is responsible for the operation of the MSPO2 plant.

The MSPO plants produce styrene monomer and propylene oxide using ethylbenzene as the raw material. Styrene monomer is used for the production of polystyrene, a plastic that is used in a wide range of products such as polystyrene foam. Propylene oxide is used for the production of propylene glycol, which is used in food, cosmetics, medicines, and other products.

Worldwide, Shell has three more plants in which styrene monomer and propylene oxide are produced by means of a process that is virtually the same as at the MSPO2 plant. Two plants are located in Singapore, at a site called Seraya, and one plant is in Nanhai, China.

In general terms, styrene monomer and propylene oxide are produced as follows: Ethylbenzene reacts with oxygen whereby it is converted into ethylbenzene hydroperoxide. The ethylbenzene hydroperoxide then reacts with propylene with the help of a catalyst[26] and is converted into propylene oxide and methylphenylcarbinol and methylphenyl ketone. The methylphenyl ketone is a by-product of this reaction. In the last step, the methylphenylcarbinol is converted into styrene monomer. The by-product methylphenyl ketone is also converted into methylphenylcarbinol in a separate process step with the help of a different catalyst. It was in this final step of the process that the accident occurred.

The explosion was in the hydrogenation Unit (4800) of the MSPO2 plant. In the reactors of Unit 4800, hydrogen is used along with a catalyst to convert methylphenyl ketone into methylphenylcarbinol. This conversion, using hydrogen, is known as hydrogenation. The reaction with hydrogen in Unit 4800 releases heat, which is dissipated by allowing liquid ethylbenzene to flow along the catalyst in the reactors. The process is called “exothermic hydrogenation reaction.” It requires a pressure increase in the reactor. Because hydrogen is very flammable when combined with the increased pressure, fire can occur in the event of a leak. This hazard places important safety requirements on the design and operation of the Unit.

In general terms, Unit 4800 consists of two reactors, two separation vessels, a combined installation with which a liquid can be heated or cooled, and an installation for condensing the gas flow. The various parts of the Unit 4800 installation are interconnected by pipes and one central pump. See Figure B.1.

Liquids and gases from the reactor are separated from each other in the separation vessels. The gases from the first separation vessel go to reactor 2, and the gases from the second separation vessel

[25] Because I do not know the name of the predecessor organization, it will be referred to by the current name, Shell Projects and Technology, in this analysis.

[26] A catalyst is a substance that influences the rate of a specific chemical reaction.

go to the flare (combustion). For, the separation vessel to function properly, it is important to achieve the correct ratio of gas and liquid. Various safety devices are used to achieve this goal.

The reactors contain a catalyst. The catalyst is used to accelerate the reaction between the substances being used in the reactors. In Unit 4800, the catalyst is in the form of cylindrical catalyst pellets. These are composed of different elements, including copper, chromium and barium. After a number of years of production, the effects of the catalyst decline and it has to be replaced. The catalyst pellets are replaced during a brief maintenance stop. The replacement of the pellets was uneventful in this case.

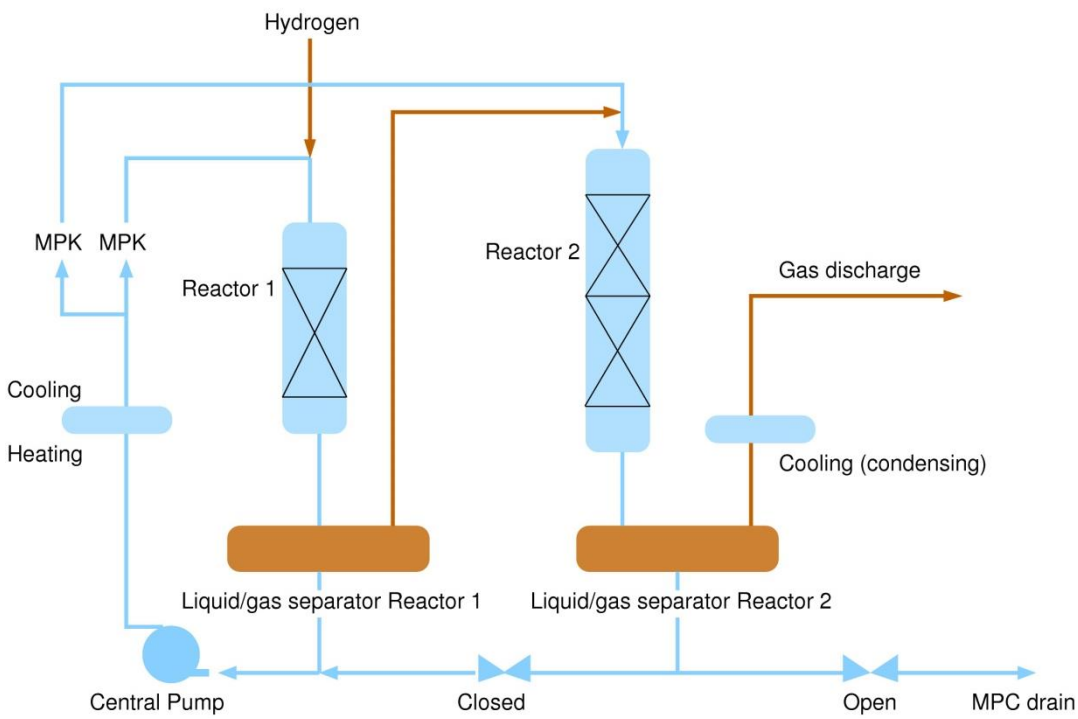

Figure B.1: Unit 4800 during normal production [taken from Dutch Safety Board report]

After the catalyst pellets have been replaced, Unit 4800 has to be restarted. This restart involves several steps: (1) release the oxygen from the Unit and then test for leaks; (2) flush the Unit with ethylbenzene to remove contamination; (3) fill the Unit with clean ethylbenzene and start circulating the ethylbenzene (called the *circulation phase*); (4) heat up the Unit (the *reheating phase*); and (5) reduce the catalyst using hydrogen (the *reduction phase*).

Circulating the ethylbenzene and heating the Unit (Steps 3 and 4) are necessary in order to wet the catalyst pellets and to raise the Unit temperature to a level that facilitates the reduction of the catalyst. The accident occurred during the reheating phase (Step 4).

Thoroughly wetting the catalyst pellets in a trickle-bed reactor[27] is critical. Wetting involves fully soaking the catalyst with ethylbenzene and keeping the pellets continuously wet. If there are localized dry zones, the heat released from a reaction cannot dissipate. The result can be an undesirable rise in the temperature of the reactors. To ensure the catalyst pellets are wet down thoroughly, enough ethylbenzene and nitrogen must be allowed to flow through the reactors and the ethylbenzene must be well distributed. These requirements are achieved by feeding ethylbenzene (liquid) and nitrogen (gas) in the correct ratios through a distribution plate in the reactors, creating a "shower effect" that distributes the liquid optimally across the catalyst pellets.

Catalyst reduction (the fifth step in restarting the reactor) can begin once the plant is at the correct temperature and hot ethylbenzene has been circulated through it for at least 6 hours. Unit 4800 never reached this step on the evening of the accident due to explosions and fire during the heating phase.

Final Summary Causal Analysis and Recommendations Resulting from the CAST Analysis

In general, there were so many flaws in the Shell Safety Management System and the behavior of almost every component of Shell Global and Shell Moerdijk that a comprehensive redesign of the Shell and Shell Moerdijk safety management system and the safety information system would seem to be appropriate. In addition, there appear to be flaws in the safety culture that should be corrected. The oversight authorities also were ineffective in preventing the accident using their procedures and legislation and a review and redesign of the oversight procedures appears to be appropriate.

Figure B.2 shows the safety control structure assumed in the CAST analysis, with flawed control and feedback contributing to the accident shown with dotted lines. As can be seen, almost the entire structure was involved.

[27] Trickle-bed reactors (used in Unit 4800) have "open" columns filled with catalyst in which a gas and a liquid flow together in the same direction under the influence of gravity,

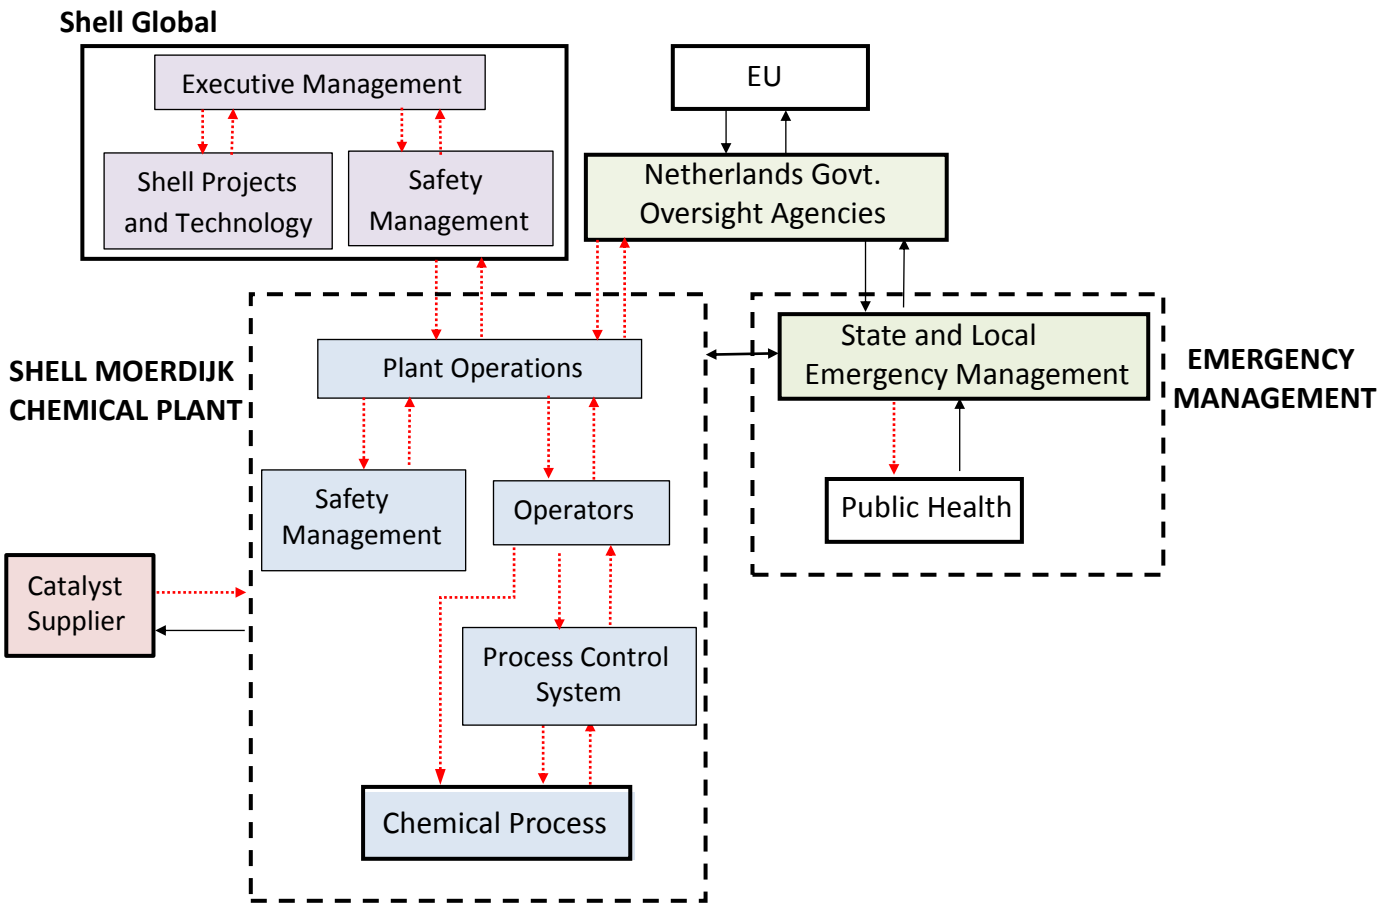

Figure B.2: Flawed Interactions in the assumed Safety Control Structure. Red dotted lines represent missing or inadequate control or feedback. Almost all the interactions were flawed in some way.

As an overview of the CAST analysis results, the following table summarizes each component's role in the accident along with the recommendations generated for that component. The reasons for the component's role in the accident would probably be augmented if the unanswered questions noted in the CAST analysis details had been included in the accident report.

Chemical Process	**Role:** None of the physical controls failed. The final physical collapse of the reactor and separation vessel after pressure reached a critical level resulted from unexpected and unhandled chemical and physical interactions. Many of these unsafe interactions were a result of design flaws in the reactor or in the safety-related controls. **Recommendations**: The physical design limitations and inadequate physical controls need to be fixed. (The potential detailed fixes are not included here; they need to be determined by a qualified chemical engineer.)
Process Control System	**Role**: The process control system was not configured to provide the necessary help to the operators during a start-up or to allow them to easily stop the process in an emergency. The reason for these design decisions rests primarily in incorrect assumptions by the designers, who thought the accident scenario was impossible. Even after previous incidents at similar plants in which these assumptions were found to be untrue, the assumptions were not questioned and revisited. **Recommendations**: The operators' knowledge and skill are most challenged during off-nominal phases, and most accidents occur during such phases and after changes are made or occur. The process control system should be redesigned to assist operators in all safety-critical, off-nominal operations (not just this restart scenario). For manual operations, the goal should be to provide all necessary assistance to the operators in decision making and in reducing attention and time pressures.
Operators	**Role**: The operators acted appropriately or at least understandably given the context, the incorrect work instructions (which they followed), and their lack of training needed to perform the work. In addition, they were provided with almost no assistance from the process control system, while many of the tasks they needed to do required intense attention, precision, mental effort, deep understanding of process dynamics, and frequent adjustments to a continually fluctuating process. The risks were not communicated properly. Management relied on the operators seeing something strange and stopping the process, but did not provide the information and training the operators needed to do this. **Recommendations**: The operators must have the appropriate skills and expertise to perform their assigned activities, and there must be someone assigned the responsibility for enforcing this requirement. A human factors study during the job analysis is needed to ensure that the operators are provided with the necessary information and work situation that allows them to make appropriate decisions under stressful conditions. Also, better automated assistance should be provided in all phases of operation, training should be provided for activities that are known to be hazardous like startup, and work instructions as well as the process for producing them need to be improved.

Plant Safety Management	**Role**: (1) The safety analysis methods used were either not appropriate, not applied or were applied incorrectly. However, the methods used complied with the Shell requirements and with the minimum required by the Dutch regulators. Safety management did not consider some technical relevant information nor investigate how ethylbenzene reacting with the catalyst could cause an explosion. Safety management at Shell Moerdijk, as is common in many places, seems to have been largely ineffectual, with lots of activity, but much of it directed to minimal compliance with government regulation. A partial explanation for their behavior is that everyone believed that a reaction between ethylbenzene and the catalyst was impossible and that the start-up process was low risk. (2) Although Shell's safety management system includes requirements for dealing with changes, the MOC procedures were not followed or implemented effectively. Risks resulting from changes made to the plant, the catalyst, the processes, and the procedures were not identified and managed. (3) The number of leaks were used as the primary leading indicator of process safety. This practice is common in the petrochemical industry. (4) Lessons from similar incidents at Nanhai and at Shell Moerdijk were not used to reduce risk. (3) Proper oversight of the generation of work instructions was not provided, which allowed unsafe work instructions to be used by the operators. **Recommendations**: While the problems specific to the explosions on 3 June 2014 should be fixed, there were a lot of weaknesses in the Shell Moerdijk safety management design and especially practices that were identified in the official Dutch Safety Agency accident report and in the CAST analysis. These need to be improved. In addition: (1) Safety management at Shell Moerdijk needs to be made more effective. Safety engineering needs to be more than just going through the motions and minimally complying with standards. (2) All work instructions should be reviewed for safety by knowledgeable people using information from the hazard analysis. [In this case, the hazard analysis was flawed too, but that is a different problem to fix.] (3) MOC procedures must be enforced and followed. When changes occur, assumptions of the past need to be re-evaluated. (4) Hazard analysis and risk assessment methods need to be improved. (5) Better leading indicators of risk need to be identified and used. (6) Procedures for incorporating and using lessons learned need to be established or improved.
Operations Management	**Role**: (1) Operations management did not identify the flaws in the risk analyses performed or the procedures used for these risk analyses. The risk analyses complied with the minimal requirements of the Dutch regulatory authorities and apparently with the Shell requirements.

	(2) Changes over time were not subjected to assessment in accordance with the MOC procedures. (3) Work instructions were created by the operators without safety engineering oversight. They did not comply with the required Shell format for such work instructions and did not include important criteria for the job such as heating rate and nitrogen flow. (4) Operations management decided to configure the process control system to control the plant during the normal production phase but not during non-production and maintenance phases. They did not think these activities were high risk and that manual operation would therefore suffice. The reasons for this decision are not in the accident report. (5) They allowed two employees from different contractors to work in the adjacent unit during the start-up, probably because they did not believe that phase was dangerous. (6) They did not assign operators to the start-up that had the qualifications required in the Safety Report. No reason is given in the accident report as to why this happened. (7) They did not ensure that lessons learned from similar plants and at Shell Moerdijk in 1999 were incorporated in the design and operation of Unit 4800. (8) They either did not follow MOC procedures or thought the procedures were unnecessary in this case. No information is in the report to determine this. (9) They conducted internal Shell Moerdijk audits that did not detect any of the clear shortcomings in practices and procedures. Not enough information is provided to determine why the audits were ineffective. **Recommendations**: (1) Establish and enforce proper MOC procedures. If changes occur, retest assumptions that could be affected by those changes. This implies that these assumptions must be recorded, leading indicators established for identifying when they may no longer be correct, and a process established for testing and responding to changes that might affect these assumptions. (2) Do a thorough review of the Shell Moerdijk SMS with emphasis on why it was unable to prevent this accident. Major factors in this accident are related to basic activities that should have been controlled by the SMS. (3) Update procedures to eliminate the causes of the accident such as lack of control and supervision of the work instruction creation and oversight processes, inadequate hazard analysis and risk assessment procedures, assignment of operators to perform the turnaround who did not have the required skills and expertise, inadequate use of lessons learned from the past, and audit procedures that did not identify the shortcomings before the accident. (4) Improve the process control system to provide appropriate assistance to operators performing functions that are outside of normal production.

Shell Projects and Technology	**Role**: The design data provided to the licensees was not usable by those creating work instructions at the plants using the technology. The design had safety-critical design flaws that were not found in hazard analyses during the initial design phase and were not fixed after receiving information about serious problems in operations at some Shell plants. These design flaws include an inadequate number of temperature sensors and the use of pressure relief valves that could not handle the pressure that occurred. Unsafe and incomplete work instructions were approved by Shell Projects and Technology for the Unit 4800 turnaround at Shell Moerdijk. Without more information about the operations at Shell Corporate, it is difficult to determine exactly why the unsafe control occurred. More questions than answers arise from the CAST analysis here, such as *Why were the design flaws introduced and how did they get through the design process? What type of hazard analysis is performed by Shell Projects and Technology (or used if it is produced by another group)? Why were identified design flaws not fixed after the incidents at Shell Moerdijk in 1999 and Nanhai in 2011? What other types of feedback (beyond incidents) is provided about the safety of their designs during operations in the Shell plants? What information about the safety aspects (hazards) of the plant design are passed from Shell Projects and Technology to the licensees of their designs? What information is included in the design book? Is the design data provided adequate for the licensees to create safe work instructions if engineers are writing the work instructions instead of operators? Did they not know who was going to be performing this task? Why did they approve unsafe work instructions that did not even follow the required Shell format? What information is provided in the Design Book specifically about start-up and the hazards of start-up? What types of hazard analysis are performed during the design process? What is the process for ensuring safety when changes are made? How are safety-related assumptions recorded and what are the triggers that initiate a re-analysis of these assumptions? What feedback do the designers receive about the operation of their designs?* **Recommendations:** Fix the design features contributing to accident. Determine how these flaws got through the design process and improve the design and design review process. Improve the design book so that it is understandable by those who are writing the work instructions and contains all the information needed to safely operate installations of the licensed technology. Improve the work instruction review process by Shell Projects and Technology to ensure the instructions are complete and safe. Review and improve the hazard analysis process used by Shell Projects and Technology (or fix it elsewhere if this group does not do their own hazard analyses).
Corporate Safety Management	**Role**: There appears to have been a flawed view of the state of risk and the effectiveness of the safety management system in Shell plants. The flawed process model is most likely related to inadequate feedback (including audits and leading indicators). Again, many questions are raised in the CAST analysis that need to be answered in order to understand the role of corporate level safety management in the accident so that more effective safety management can be provided in the future. Almost nothing about safety management at the Corporate level is included in the accident report.

	Recommendations: Improve Shell safety audits. Review all risk assessment and hazard analysis processes and, in general, improve their approach to both safety analysis and safety management. Shell is not alone among the large oil companies in needing to update their methods. The petrochemical industry has too many accidents and incidents that are avoidable. More specifically, the accident report says that Shell should "evaluate how risk analyses are performed and make changes. This should include procedures and policies requiring re-evaluation of earlier presumptions and assumptions. Conduct new risk analyses, put adequate measures in place and ensure that the team that performs these analyses has sufficient critical ability. Pay particular attention to assumptions based on risks that had previously been ruled out." Evaluate and improve the corporate safety management system. Improve procedures for learning from process safety-related incidents. Create better feedback mechanisms (including audits and leading indicators) and procedures for learning from incidents.
Executive-Level Corporate Management	**Role**: Corporate management is responsible to ensure that an effective safety management system is created. Clearly typical policies of an effective safety management system were violated at both Shell Corporate and Shell Moerdijk. The group overseeing safety at the Shell corporate level was not effective. There is nothing included in the accident report about the assigned safety-related responsibilities for Corporate management. The Baker Panel report on the Texas City explosion found that BP corporate management did not have assigned responsibilities for safety, which instead was treated as a local responsibility. This abdication of responsibility (a practice promoted by HRO, which BP follows) was identified as a major contributor to the Texas City explosion [Baker 2007]. Is this a problem in general in the petrochemical industry? There is also nothing included about context in the accident report that might explain why standard executive-level responsibilities for safety were not effectively carried out. There seems, however, to be a safety culture problem at Shell. Is this a problem in the Oil and Gas industry as a whole? The accident report notes that the Safety Management System was integrated with the Business Management System at Shell Moerdijk. Was this also true at the corporate level? This is a very poor practice (and was a factor in the Deepwater Horizon accident). Safety risk assessments need to be kept separate from business risk assessments so that information is not hidden from high-level decision-makers. **Recommendations**: Review the SMS design and determine why it did not prevent obvious violations of policy such as shortcomings in safety studies, management of change, learning from accidents, not following regulations (e.g., having experienced operators and following the format for work instructions). Determine why audits were not effective in finding such obvious violations of procedures. While it is possible that this was the first time such lapses have occurred, it is highly unlikely. Strengthen audit procedures, including identifying better leading indicators of

	increasing risk than simply the number of leaks and create other forms of feedback to identify when the safety management system is drifting off course and risk is increasing. Establish better feedback channels to ensure that management of change procedures and corporate safety policy are being followed.

Catalyst Manufacturer	**Role**: The changes made in the catalyst were not pointed out to Shell, but they were included in a new safety information sheet. While the catalyst manufacturer cannot determine the impact of their changes on a customer, there should be some clear alert, other than simply changing information in a document, that changes have been made so that the customers are aware of them. **Recommendations**: Change contractual relationships between Shell and its suppliers to ensure that potentially critical changes are communicated. Make changes within information sheets so they are clear and obvious.

Dutch Regulators	**Role**: The accident report implies that regulators gave Shell Moerdijk a pass on behavior that might have been labeled violations. Plant scenario deficiencies should have been considered a violation but were not. Scenarios were not up to date or were incomplete. Working under limited resources and time is difficult under any supervision model but system-level supervision has major limitations in ensuring public safety. The accident investigation showed many flaws in Shell Moerdijk operations safety management as defined and as implemented. So what is wrong with the supervision model that the regulators did not detect them? **Recommendations**: Better supervision of the highest risk activities is needed, including turnarounds. Regulators need to oversee and ensure that strict procedures are being used for the most dangerous activities and that the safety management system is operating effectively and following its own rules. Operating under limited resources does not preclude doing something effective; it simply requires a more intelligent selection of activities that are performed. There is a need for better evaluation procedures and oversight of safety management system effectiveness. The regulators should rethink system-level supervision to ensure that they provide effective oversight in preventing accidents like the Shell Moerdijk explosions and fire.
Emergency Services	**Role**: Emergency services were mostly effective in carrying out their responsibilities, but some deficiencies, particularly in communication, were uncovered in the accident response. **Recommendations**: Several deficiencies were uncovered in LCMS and NL-Alert communication protocols during this incident. While they did not lead to loss of life because of the nature of this accident, they could under other circumstances and what was learned from this case should be used to improve the system, including why many people used WhatsApp instead of LCMS and how the official system can incorporate those features. Accidents create an opportunity to look closely at the

	actual operation of our systems in times of stress and provide a learning opportunity that should not be wasted.

For the factors that spanned the entire control structure, not much information required for the CAST analysis was provided in the accident report. Some weaknesses are implied by what is included and some general recommendations can be derived.

Safety Management System	Evidence of an overall inadequate safety control system and therefore safety management system in the report includes: unsafe situations were overlooked, internal procedures were not properly followed, lessons were not learned from previous incidents, incorrect assumptions about basic chemical reactions were not re-evaluated after evidence surfaced that they were incorrect, changes were not managed and controlled, inadequate hazard analysis and risk assessment procedures were used, recommendations from previous incidents and accidents were never implemented, and oversight of critical activities was missing. In summary, the Safety Management System at Shell Moerdijk did not prevent unsafe situations from being overlooked or internal procedures from not being followed. There is no information in the accident report about who created the SMS or who was responsible for ensuring that it was working properly. **Recommendations**: The design of the entire safety management system should be evaluated and improved. In addition, the integration of the safety management system and the business management system should be carefully examined to ensure that hazard and risk-related information is being effectively communicated to decision makers and not stopped at inappropriately low levels.
Safety Information System	No information is provided about the safety information system but it appears that people were making decisions without having appropriate information. **Recommendations**: The safety information system is so critical to the achievement of high safety that Shell and Shell Moerdijk should evaluate the existing system and perhaps redesign it.
Safety Culture	The accident report did not cover the safety culture in depth, but what is included seems to point to what has been labeled as a “compliance culture,” where the bulk of the effort is simply complying with standards and the requests of regulators and not proactively taking steps to improve safety. The accident report points to unsafe behavior that seems to imply safety was not a high priority, such as overlooking unsafe behavior and warnings, not adhering to internal procedures, not making changes after previous incidents, and not evaluating assumptions after changes occur. The Hearts and Minds Safety Culture Program used by Shell has serious weaknesses. The “culture ladder” is vaguely defined (“we do a lot every time we have accidents”). Strangely, the company seems to be satisfied with their self-assessed current level in this program of *Calculative* (which is described in the

	accident report as the “required” level), but even that level does not seem to have been achieved(nor even the so-called “lower” levels). **Recommendations**: Shell Moerdijk and Shell Corporate should do a thorough study of their safety culture and determine why it was not strong enough to prevent the events in 2014.
Communication and Coordination	**Recommendations**: Communication channels, especially feedback channels, should be examined to determine whether they are effectively conveying the information necessary to operate safely. An important part of this includes performing a human factors analysis of the information provided to the operators and the potential for human errors created by the design of the process and particularly by the design of the process control system.
Management of Change	**Role**: A large number of both planned and unplanned changes contributing to the accident were not assessed for risk. **Recommendations**: The Management of Change procedures should be evaluated to determine why they were not effective in this case and appropriate improvements implemented.

Appendix C: The Bad Apple Theory of Accident Causation

Blaming workers in accident investigation is based partially on the outdated and erroneous "bad apple" theory.[28] This concept says that there are a certain number of "bad apples" or a few people in a system that are responsible for most of the accidents. Most everything would be safe, in this view, if it were not for a few unreliable humans.

The concept goes back to 1925 when both German and British psychologists were convinced they would solve the safety problem by identifying and getting rid of the bad apples in an organization. They used statistical analysis over 50 years to determine that there were a cohort of "accident-prone" workers. These were people with personal characteristics that, they claimed, predisposed them to making errors and thus precipitating accidents. The data seemed to imply that a small percentage of people is responsible for a large percentage of accidents. If those people were removed, the system would become much safer.

The bad apple theory existed until WW II, when the complexity of the systems we were creating and that humans had to work within, started to increase dramatically. The theory was finally put to rest in 1951 by two statisticians, Arbous and Kerrich.

The theory did not work because of a major statistical flaw in the argument. For the accident-prone or bad apple theory to work, the risk of error and accidents must be equal across every system. But, of course, it is not. Newer views of accident causation conclude that personal characteristics do not carry as much explanatory power for why accidents occur as context does.

> *'When faced with a human error problem you may be tempted to ask 'Why didn't they watch out better? How could they not have noticed?' You think you can solve your human error problem by telling people to be more careful, by reprimanding the miscreants, by issuing a new rule or procedure. They are all expressions of the 'Bad Apple Theory' where you believe your system is basically safe if it were not for those few unreliable people in it. This old view of human error is increasingly outdated and will lead you nowhere.'* – Sidney Dekker*, A Field Guide to Understanding Human Error*, 2002.

The bad apple theory is especially prevalent in medicine, but it permeates almost every industry in the form of argument that most accidents are caused by human operators. If, for example, one posits that 5% of bad doctors cause most accidents, then simply identifying and getting rid of the doctors who get the most complaints and are involved in adverse events should drastically reduce medical error. Unfortunately, it does not. It may simply get rid of a group of doctors who do the really difficult, tricky work (e.g., some oncological cases with a negative prognosis).

More generally, if there are system features that are likely to induce human mistakes under some circumstances (e.g., poor facilities with limited resources and financial and other stresses), then swapping the humans involved, but not changing the system features that are creating the erroneous behavior, will have little impact on the accident rate. We know now that the design of systems has the potential to create various classes of errors such as mode confusion or, conversely, to reduce them. Surprisingly, most humans don't realize they are being influenced by their environment. Think about how you are fooled by optical illusions, even when you know they are there.

---

[28] Much of the information in this appendix comes from Sidney Dekker. We wrote an editorial together on the systems approach to medicine (Dekker, S. and Leveson, N. The systems approach to medicine: Controversy and misconceptions, *British Medical Journal: Quality and Safety*, 2014 and "The bad apple theory won't work", British Medical Journal: Quality and Safety, Nov. 2014). Dekker also wrote about the "new view" of human error and the "bad apple" theory in his 2002 book *The Field Guide to Understanding Human Error,* Ashgate Publishers, 2006.

As a simple physical example, consider the designs for a stove top shown in Figure C.1 from Don Norman's *The Design of Everyday Things* (Basic Books, 2013). These designs, which are common even today, require labels to connect the knobs with the cooking surfaces. Even after fairly extensive use, one may still need to read the labels to make the proper association or, if in a hurry, operate the wrong knob.

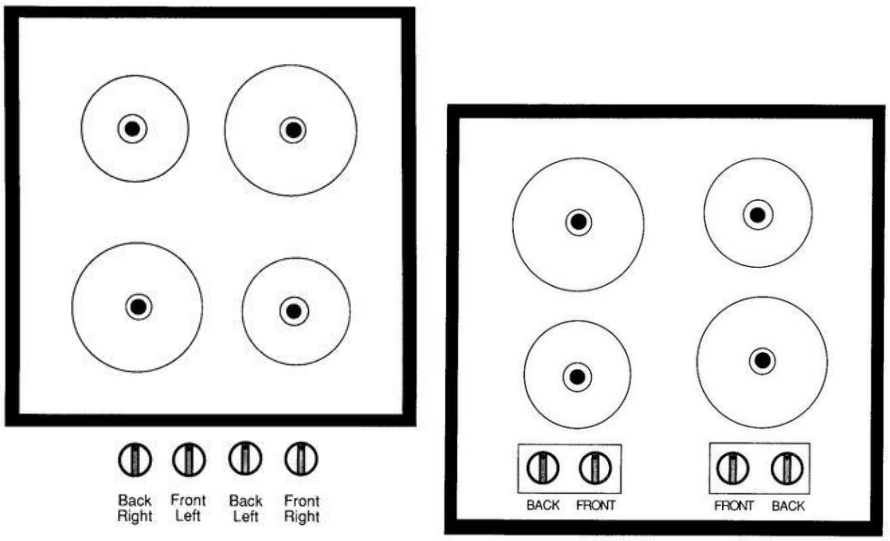

Figure C.1: Two designs of an error-prone stove top (Adapted from Don Norman, The Design of Everyday Things, Basic Books, 2013).

The designs in Figure C.2 eliminates most of these errors and, in fact, the need for labels is eliminated.

The examples of physical designs that induce human errors (or at least do not prevent them) shown in Figure C.1 are relatively simple to avoid in system design (which does not explain why so many stove tops today still use the error-prone design). But the problem is greatly exacerbated when cognitively complex decision making is required by operators. Most of our systems today require humans to work with computers and other complex designs that can more easily create an incorrect mental model on the part of the human about the state of the system or the state of the automation. Blaming the human, then, for erroneous behavior without looking at why that behavior occurred does nothing to reduce accidents due to the flawed system design. Pilots today, for example, are doing the wrong thing in a certain situation because of factors like inconsistent behavior of the automation, misunderstanding about how the automation is designed, confusion about the current mode of the aircraft or its automation, etc. Identifying these error-inducing design features during accident causal analysis, rather than simply stopping with "operator error," will allow us to improve safety in both current systems and in our future designs.

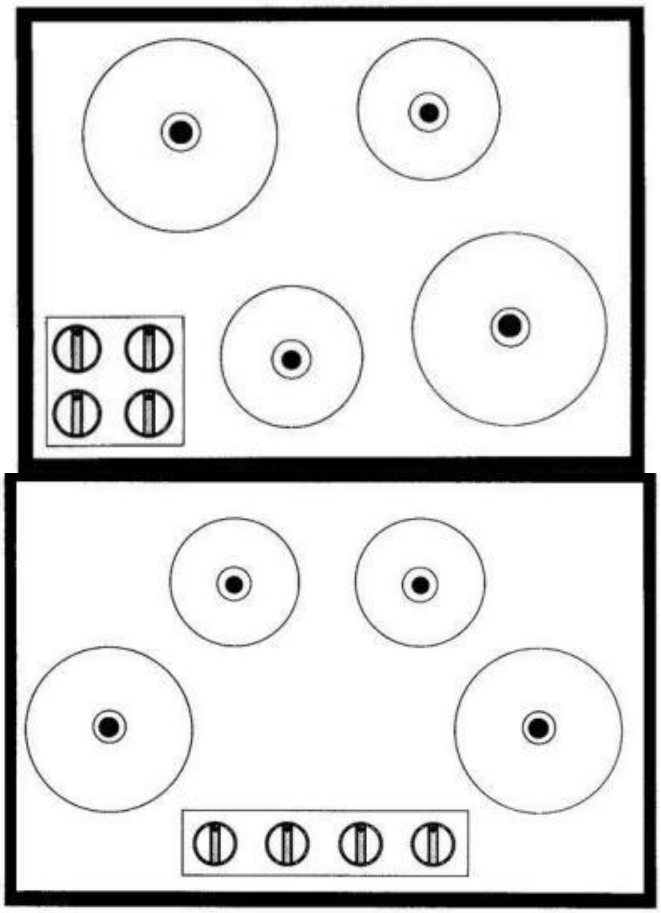

Figure C.2: Less error-prone designs (Adapted from Don Norman, The Design of Everything Things, Basic Books, 2013).

Appendix D: Factors to Consider When Evaluating the Role of the Safety Control Structure in the Loss

The list below (from *Engineering a Safer World*, Chapter 13) can be used in identifying inadequate controls and control structures when performing incident and accident analysis. It is not meant to be exhaustive and may need to be supplemented for specific industries and safety programs.

This list contains general responsibilities and does not indicate how they should be assigned. Appropriate assignment of the responsibilities to specific people and places in the organization will depend on the management structure of each organization. Each general responsibility may be separated into multiple individual responsibilities and assigned throughout the safety control structure, with one group actually implementing the responsibilities and others above them supervising, leading or directing, or overseeing the activity. Of course, each responsibility assumes the need for associated authority and accountability, as well as the controls, feedback, and communication channels necessary to implement the responsibility.

General Management

- Provide leadership, oversight, and management of safety at all levels of the organization.
- Create a corporate or organizational safety policy. Establish criteria for evaluating safety-critical decisions and implementing safety controls. Establish distribution channels for the policy. Establish feedback channels to determine whether employees understand it, are following it, and whether it is effective. Update the policy as needed.
- Establish corporate or organizational safety standards and then implement, update, and enforce them. Set minimum requirements for safety engineering in development and operations and oversee the implementation of those requirements, including any contractor activities. Set minimum physical and operational standards for hazardous operations.
- Establish incident and accident investigation standards and ensure recommendations are implemented and effective. Use feedback to improve the standards.
- Establish management of change requirements for evaluating all changes for their impact on safety, including changes in the safety control structure. Audit the safety control structure for unplanned changes and migration toward states of higher risk.
- Create and monitor the organizational safety control structure. Assign responsibility, authority, and accountability for safety.
- Establish working groups.
- Establish robust and reliable communication channels to ensure accurate management risk awareness of the development system design and the state of the operating process. These channels should include contractor activities.
- Provide physical and personnel resources for safety-related activities. Ensure that those performing safety-critical activities have the appropriate skills, knowledge, and physical resources.
- Create an easy-to-use problem reporting system and then monitor it for needed changes and improvements.
- Establish safety education and training for all employees and establish feedback channels to determine whether it is effective along with processes for continual improvement. The education should include reminders of past accidents and causes and input from lessons learned and trouble reports. Assessment of effectiveness may include information obtained from knowledge assessments during audits.

- Establish organizational and management structures to ensure that safety-related technical decision making is independent from programmatic considerations, including cost and schedule.
- Establish defined, transparent, and explicit resolution procedures for conflicts between safety-related technical decisions and programmatic considerations. Ensure that the conflict resolution procedures are being used and are effective.
- Ensure that managers who are making safety-related decisions are fully informed and skilled. Establish mechanisms to allow and encourage all employees (including front-line operators) and contractors to contribute to safety-related decision making.
- Establish an assessment and improvement process for safety-related decision making.
- Create and update the organizational safety information system.
- Create and update safety management plans.
- Establish communication channels, resolution processes, and adjudication procedures for employees and contractors to surface complaints and concerns about the safety of the system or parts of the safety control structure that are not functioning appropriately. Evaluate the need for anonymity in reporting concerns.

Development

- Implement special training for developers and development managers in safety-guided design and other necessary skills. Update this training as events occur and more is learned from experience. Create feedback, assessment, and improvement processes for the training.
- Create and maintain the hazard log. Establish and maintain documentation and tracking of hazards and their status.
- Establish working groups.
- Design safety into the system using system hazards and safety constraints. Iterate and refine the design and the safety constraints as the design process proceeds. Ensure the system design includes consideration of how to eliminate or reduce contextual factors that cause or contribute to unsafe operator behavior that, in turn, contributes to system hazards. Distraction, fatigue, etc. are risk factors resulting from design that is dependent on humans performing in a way the designer imagined they would rather than behaving as normal humans would in such situations.
- Document operational assumptions, safety constraints, safety-related design features, operating assumptions, safety-related operational limitations, training and operating instructions, audits and performance assessment requirements, operational procedures, and safety verification and analysis results. Document both what and why, including tracing between safety constraints and the design features to enforce them.
- Perform high-quality and comprehensive hazard analyses and make these available and usable when safety-related decisions need to be made, starting with early decision making and continuing through the system's life. Ensure that the hazard analysis results are communicated in a timely manner to those who need them. Establish a communication structure that allows communication downward, upward, and sideways (i.e., among those building subsystems). Ensure that hazard analyses are updated as the design evolves and test experience is acquired.
- Train engineers and managers to use the results of hazard analyses in their decision making.
- Maintain and use hazard logs and hazard analyses as experience is acquired. Ensure communication of safety-related requirements and constraints to everyone involved in development.
- Gather lessons learned in operations (including accident and incident reports) and use them to improve the development processes. Use operating experience to identify flaws in the development of safety controls and implement improvements.

Operations

- Create an operations safety management plan
- Develop special training for operators and operations management to create needed skills and update this training as events occur and more is learned from experience. Create feedback, assessment, and improvement processes for this training. Train employees to perform their jobs safely, understand proper use of safety equipment, and respond appropriately in an emergency.
- Establish working groups.
- Maintain and use hazard logs and hazard analyses during operations as experience is acquired.
- Ensure all emergency equipment and safety devices are operable at all times during hazardous operations. Before safety-critical, non-routine, potentially hazardous operations are started, inspect all safety equipment to ensure it is operational, including the testing of alarms.
- Perform an in-depth investigation of any operational anomalies, including hazardous conditions (such as water in a tank that will contain chemicals that react to water) or events. Determine why they occurred before any potentially dangerous operations are started or restarted. Provide the training necessary to do this type of investigation and create proper feedback channels to management.
- Create management of change procedures and ensure they are being followed. These procedures should include hazard analyses on all proposed changes and approval of all changes related to safety-critical operations. Create and enforce policies about disabling safety-critical equipment.
- Perform safety audits, performance assessments, and inspections using the hazard analysis results as the preconditions for operations and maintenance. Collect data to ensure safety policies and procedures are being followed and that education and training about safety is effective. Establish feedback channels for leading indicators of increasing risk.
- Use the hazard analysis and documentation created during development and passed to operations to identify leading indicators of migration toward states of higher risk. Establish feedback channels to detect the leading indicators and respond appropriately.
- Establish communication channels from operations to development to pass back information about operational experience.
- Perform in-depth incident and accident investigations, including all systemic factors. Assign responsibility for implementing all recommendations. Follow up to determine whether recommendations were fully implemented and effective.
- Perform independent checks of safety-critical activities to ensure they have been done properly.
- Prioritize maintenance for identified safety-critical items. Enforce maintenance schedules.
- Create and enforce policies about disabling safety-critical equipment and making changes to the physical system.
- Create and execute special procedures for the startup of operations in a previously shutdown unit or after maintenance activities.
- Investigate and reduce the frequency of spurious alarms.
- Clearly mark malfunctioning alarms and gauges. In general, establish procedures for communicating information about all current malfunctioning equipment to operators and ensure they are being followed. Eliminate all barriers to reporting malfunctioning equipment.
- Define and communicate safe operating limits for all safety-critical equipment and alarm procedures. Ensure that operators are aware of these limits. Assure that operators are rewarded for working within the limits and following emergency procedures, even when it turns out no emergency existed. Provide for tuning the operating limits and alarm procedures over time as required.
- Ensure that spare safety-critical items are in stock or can be acquired quickly.

- Establish communication channels to plant management about all events and activities that are safety-related. Ensure management has the information and risk awareness they need to make safe decisions about operations.
- Ensure emergency equipment and response services are available and operable to treat injured workers.
- Establish communication channels to the community to provide information about hazards and necessary contingency actions and emergency response requirements.

Appendix E: Basic Engineering and Control Concepts for Non-Engineers

In reviews of an early draft of our STPA handbook, we discovered that we had assumed some engineering background that not all readers had: Users of STPA and CAST are not necessarily trained in engineering. This appendix provides an introduction to basic engineering concepts that may be unfamiliar to some users of this handbook.

I have attempted to write the sections in this appendix to be independent so that readers can pick and choose sections that interest them. No assumption is made about any particular educational background in covering these topics, which include:

- What is a "system"?
- Basic systems engineering concepts and terminology
- The concept of "control" in engineering
- Systems theory vs. complexity theory

What is a "System"?

The most basic concept underlying everything else in this handbook is that of a *system*.

> <u>*System*</u>: A set of things (referred to as system components) that act together as a whole to achieve some common goal, objective, or end.

Some definitions of a system leave out the "goal" or "objective" and state essentially that a system is a connected set of things or parts forming a unified whole. This definition does not capture the way the term "system" is usually used. A shoe, a star, and a locomotive are a set of things or potential parts of a system, but most people would not normally consider this to be a system. It could be considered as a system if a purpose could be conceived for considering these individual things together; the purpose is basic to the concept.

Other definitions state that the components of the system must be interdependent or connected or interacting, but none of these conditions are really necessary to have a system. They also constrain the definition in a way that excludes things that usually are considered to be systems. The system components may be either directly or indirectly connected to each other, with the latter including connections involving the system purpose only and various types of indirect interdependencies.

The consequence of this definition is that a goal or objective for the system is fundamental. But there may be different objectives for those defining and viewing systems. An airport is used throughout this section to illustrate various points. To a traveler, the purpose of an airport may be to provide air transportation to other locales. To local or state government, an airport may be a means to increase government revenue and economic activity in the area of the airport. To the airlines, the purpose of an airport may be to take on and discharge passengers and cargo. To the businesses at the airport, the purpose is to attract customers to whom they can sell products and services. Therefore, when talking about a system, it is always necessary to specify the purpose of the system that is being considered.

> *A system is an abstraction, that is, a model conceived by the viewer.*

The observer may see a different system purpose than the designer or focus on different relevant properties. Specifications, which include the purpose of the system, are critical in system engineering. They ensure consistency of mental models among those designing, using, or viewing a system, and they enhance communication. Notice that different components of the airport may be included in a

particular “airport” system, such as passenger check-in counters, ramps to the planes, and taxiways in the airline view of an airport versus shops and customers from the commercial view. Note again that these are models or abstractions laid upon the actual physical world by human minds, i.e., a system lies in the eye of the beholder. The components that are considered in any “airport system” or subsystem and the role they play in the system as a whole may be different for each concept (model) of an airport system or of airport subsystems. The basic idea here is that the purpose or goals of the system being considered must be specified and agreed upon by those modeling and analyzing a particular system and that these aspects of systems are abstractions or models imposed by the viewer on the real-world objects.

For engineered (man-made) systems, the purpose is usually specified before creating the system although observers of the system may interpret the purpose differently than originally intended by the designers. For natural systems, behavior may be interpreted as purposeful by the observers of the system.

There are some basic assumptions underlying the concept of a system:

- The system goals can be defined
- Systems are atomistic, that is, they can be separated into components with interactive relationships between the components regardless of whether the interactions are direct or indirect.

Systems themselves may be part of a larger system or be divisible into subsystems (viewed as components of the overall system). Note that the definition here is what is called “recursive” in mathematics and computer science. That is, the definition is made in terms of itself. A system is usually part of a larger system and it can be divisible into subsystems. Figure E.1 shows a typical abstraction hierarchy for a system labeled A (for example the system “airport”) and for three subsystems, which in turn can be conceived as systems themselves.

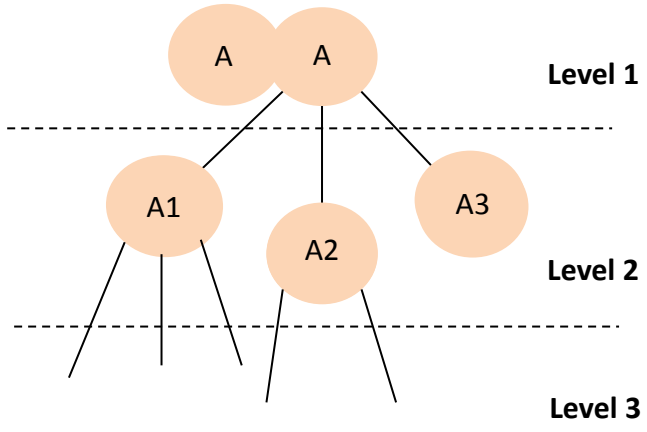

Figure E.1: The abstraction System A may be viewed as composed of three subsystems. Each subsystem is itself a system.

Figure E.1 is not a connection diagram. Rather it is a hierarchical abstraction that shows the different views of a system at different levels of abstraction or modeling. Level 1 conceives of the system as a whole. Level 2 visualizes it as being composed of three subsystems. Figure E.2 shows an abstraction where A itself is conceived as part of a larger system AB.

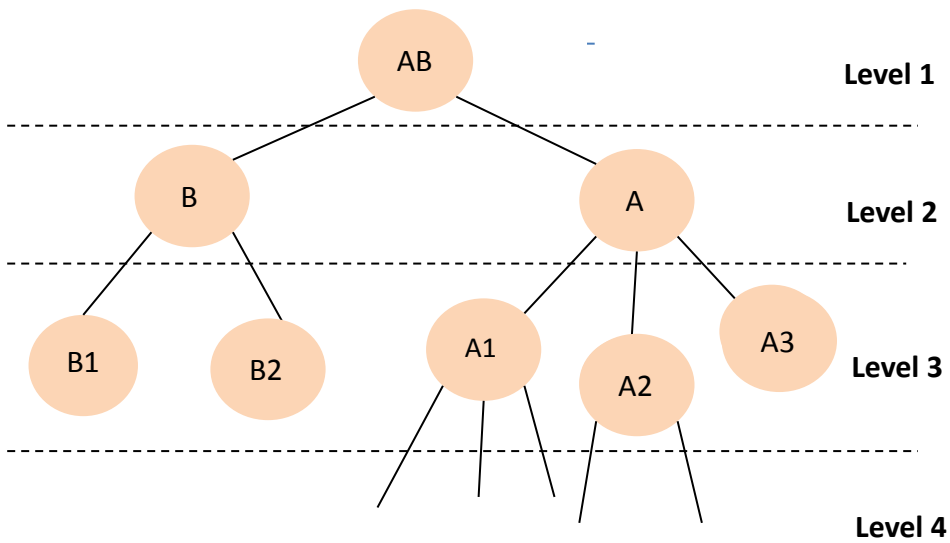

Figure E.2: System A can be viewed as a component (subsystem) of a larger system AB

The recursive nature of the definition of a system is important because many people have suggested that "systems of systems" must be treated differently than systems. In fact, the same general system engineering and analysis methods and techniques are applicable to all systems. A "system of systems" is just a "system" when using the definition of a system as defined in system theory and as the concept is used in engineering. When one puts together several systems, you get not a system of systems but an entirely new and unique system. It is not a summation of several systems.

Systems have *states*. A state is a set of relevant properties describing the system at any time. Some properties of the state of an airport viewed as an air transportation system may be the number of passengers at a particular gate, where the aircraft are located and what they are doing (loading passengers, taxiing, taking off, or landing). The components of the state that are relevant depend on how the boundaries are drawn between the system and its environment.

The *environment* is usually defined as the set of components (and their properties) that are not part of the system but whose behavior can affect the system state. Therefore, the system has a state at a particular time and the environment has a state. The concept of an environment implies that there is a boundary between the system and its environment. Again, this concept is an abstraction created by the viewer of the system and need not be a physical boundary. What is part of the system or part of the environment will depend on the particular system and its use at the time.

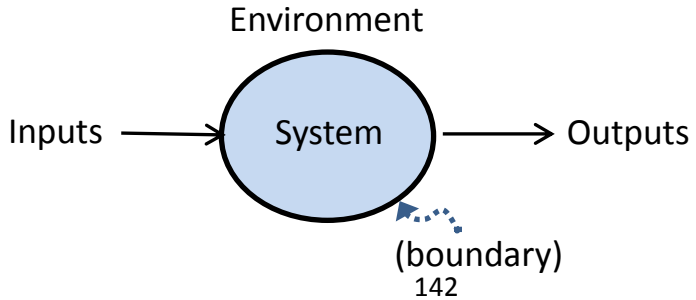

System *inputs* and *outputs* cross the system boundary. This concept is usually called an *open* system. There is also a concept of a *closed system* (which has no inputs or outputs), but it does not have much relevance for the engineered systems with which we are most concerned in system safety.

The Concept of "Control" in Engineering

The concept of control is important in the conception, design, and operation of a system. A system can conceivably operate without any type of control over its behavior, but it would be difficult to ensure that the system satisfied its goals while at the same time constraining the way it achieves those goals so that adverse, unwanted consequences do not also result.

Control is basic to engineering, especially complex systems. Engineers are dealing not with living things, but with man-made things —automobiles, dams, aircraft, chemical plants, spacecraft, washing machines, robots—that are built to accomplish some goal. If they are not inert (like a bridge), there is a need to operate or control their operation in terms of the states that they get into.

Non-engineers sometimes interpret the word "control" differently than the way it is used in engineering. Control is not about forcing individual compliance to specified roles and procedures through some type of militaristic-style authority. Rather, control is being used here in a very broad sense and may be implemented through design controls like interlocks and fail-safe design, through process controls such as maintenance and manufacturing procedures, or even social constructs such as culture, incentive structures, and individual self-interest. Without the imposition of some form of control, there would be no way to ensure that the system goals and constraints are achieved.

Simple design features acting as controls often can improve safety without also increasing complexity and cost, but this usually requires considering safety early in the design process. The most effective way to increase safety through design is to use a design that eliminates hazardous states. Next best is to design such that hazards are rare and easy to handle if they do occur. A design that reacts to the occurrence of hazards by minimizing the resulting damage is clearly the least desirable form of hazard control but often such controls must be included.

A control loop is thus only one form of control. I emphasized it in *Engineering a Safer World* because the other aspects of basic design for safety were covered so extensively in my previous book titled *Safeware*. In addition, system theory emphasizes the concept of the control loop or active control. It is not possible to teach the techniques used to design for safety in this appendix, but some basic concepts such as the distinction between p*assive* and *active* control may be helpful in understanding general safety design techniques.

There are three types of design techniques to "control" safety. The first is to use an intrinsically or inherently safe design. An intrinsically safe design is one that is not capable of producing the hazard of concern. For example, the system design may not have sufficient energy to cause an explosion. If the concern is with the use of potential toxic substances, there may be non-toxic substitutes. If the potential exists for lost information over a communication network, then more direct communication may be used if it is feasible.

If an intrinsically safe design is not possible, the use of passive control mechanisms may be possible. Passive controls do not require a positive action in order to be effective, but instead rely on basic physical principles, such as gravity. A physical interlock that prevents the system from entering a hazardous state without any overt control action is a passive control mechanism. Examples include:

- A pressure sensitive mat or light curtain that shuts off power to a robot or dangerous machinery if someone comes near.
- A physical device that ensures the correct sequencing of valve turn-off and turn-on or ensures that two valves cannot both be on or off at the same time.
- A freeze plug in a car's engine cooling system where expansion will force the plug out rather than crack the cylinder if water in the block freezes.
- A fusible plug in a boiler that becomes exposed if there is excessive heat and the water drops below a predetermined level. In that event, the plug melts and the opening allows the steam to escape, which reduces the pressure in the boiler and prevents an explosion.
- A speed governor that does not allow speeds over a specific limit or relief valves that maintain pressure below dangerous levels.

Passive controls are often designed so that the system essentially fails into a safe state. In fact, the basic approach to aircraft safety is to build the system to be fail safe. Some examples of fail-safe design are:

- Old railway semaphores that used weights to ensure that if the cable (controlling the semaphore) broke, the arm would automatically drop into the STOP position.
- Air brakes that are held in an off position by air pressure. If the brake line breaks, air pressure is lost and the brakes are automatically applied.
- A deadman switch where the operator applies continual pressure. If something happens to the operator and the pressure is not applied, the system fails into a safe state.
- A bathyscope where ballast is held in place by magnets. If electrical power is lost, the ballast is released and the bathyscope ascends to the surface.
- Designs where the effects of a failure (such as fan blades being ejected when an aircraft engine fails) are contained so that adjacent components are not affected.

An *active control*, in contrast, requires a hazardous condition to be detected and corrected. Active controls today often involve the use of computers and the basic control loops described and used throughout this handbook. Notice that there is more that can go wrong with active controls, e.g., the failure or hazardous state may not be detected or it may be detected but not corrected or it may not be corrected in time to prevent a loss. Passive controls are more reliable as they usually depend on physical principles and simpler designs while active control depend on less reliable detection and recovery mechanisms and usually much more complex designs. Passive controls, however, tend to be more restrictive in terms of design freedom and are not always feasible to implement in complex systems.

Active controls usually involve a control loop.

What is a Control Loop?

A control system commands, directs, or regulates the behavior of other devices or processes using control loops. Such control systems may range from a single home heating controller using a thermostat to control a boiler to large industrial control systems used for controlling complex processes and machinery.

A very basic control loop is shown below (and elsewhere in this handbook):

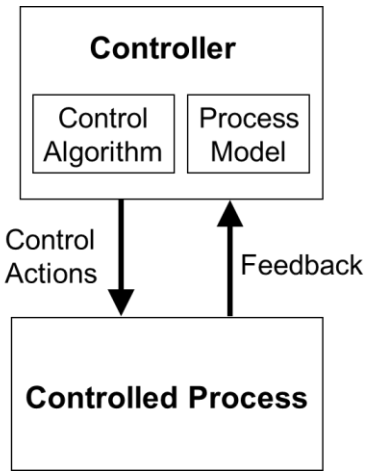

In order to control a process, the controller must have a *goal* or goals, which can include maintaining constraints on the behavior of the controlled process. In addition, the controller must have some way to affect the behavior of the controlled process, i.e., the state of the controlled process or system. These are labeled control actions in the figure. In engineering, control actions are implemented by *actuators*. In order to know what control actions are necessary, the controller must be able to determine the current state of the system. In engineering terminology, information about the state of the system is provided by *sensors* and is called *feedback*. Finally, the controller must contain some model of the state of the controlled process as explained elsewhere in the handbook.

As an example, a simple thermostat may have a goal of maintaining a temperature within a specific range. Feedback provides information to the controller about the current state of the controlled process, in this case the temperature of the air in the room. The feedback is used to update the controller's model of the process. Other information may also be part of this process model, such as environmental information (outside air temperature) or basic information about natural temperature fluctuations. The controller uses the process model to decide what control actions to provide, e.g., apply hot air or cool air, in order to change the state of the room temperature (controlled process) to stay within the required range.

There are two general operational modes for a control loop: feedback control and feedforward control. *Feedback control* is what was described in the paragraph above. When driving, a driver may read the speedometer (feedback) and decide to brake or step on the accelerator to keep the automobile's speed at a desired level.

An example of *feedforward control* occurs when the driver approaches a hill. The driver could wait until the car starts to slow down and then step on the accelerator, but more likely an experienced driver will anticipate that acceleration will be required to maintain a constant speed when going up the hill and the driver will increase acceleration before the car actually slows down. Thus, in feedforward control, the controller uses a model of the current state of the process (in this case the car speed) and the future (operating on an incline) and then provides a control action without specific feedback to identify the need.

Often feedback and feedforward control are used in tandem. In our example, the driver might anticipate the need to step on the accelerator before feedback is received using feedforward control, but feedback can be used to adjust the feedforward control actions so that a constant speed is maintained on the incline if the feedforward calculations are not perfect.

The engineering concept of control (which is based on system theory) can also be applied to social and organizational systems and has been used extensively in management theory and social sciences.

What is Systems Engineering?

System engineering is the attempt to put structure into the design and construction of a system in order to improve the results of the engineering effort. For relatively simple systems, system engineering may not be necessary. Systems engineering came into being after World War II when we started building significantly more complex systems, especially missile and other defense systems, and found that an informal design and construction process often resulted in systems that did not satisfy the goals envisioned for the system and the constraints on its potential behavior, i.e., constraints how it achieved the goals. Lack of a structured process also led to projects being late and over budget.

System engineering is a large subject, but only two concepts are necessary to understand this handbook. The first is the concept of the system engineering phases and process. The most basic model of these phases is shown in Figure E.3.

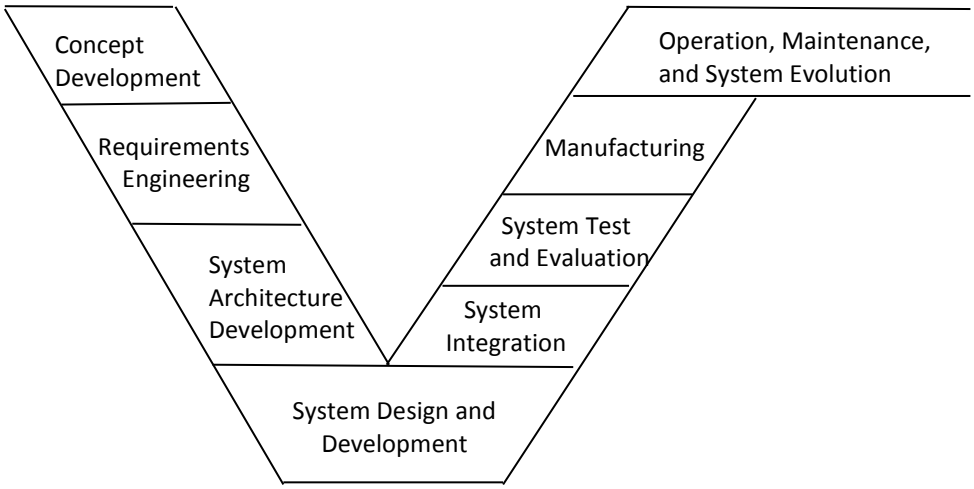

Figure E.3: The basic system engineering "V" model

This model is called the V-model, although different labels may be placed on the various components. Sometimes the phases are drawn in a straight line (if the line tilts downward it is called the "waterfall" model), but there is no difference except in the way it is drawn. The basic idea is that the best results ensue if a structured process is used, where the system concepts are first developed and the basic goals and constraints on the system are defined. There may be various feasibility and other analyses done at this early time in development to stop the process from going down paths that require retracing later. Requirements engineering is usually the next step in the process where detailed requirements are developed for the system given the basic goals and constraints earlier agreed upon. System architecture development means that a basic architecture or high-level design of the system is created before detailed design and development. The left side of the "V" represents the basic design process.

The right side of the V shows the later stages in development when assurance and manufacturing take place. Sometimes operation of the system is included (which is then called a life-cycle model)—as

shown in Figure F.3. On the right side of the V-model process, the various components that have been designed and constructed are integrated and go through a testing and evaluation phase. They are then manufactured and used.

The figure above is purposely simplified to reduce clutter and only reflects the overall process components. There is always lots of bouncing around the phases and not one straight-line process. Requirements are never totally specified at the beginning; the design process itself will generate more understanding of the problem being solved and therefore result in adding to or changing the original requirements and constraints. Serious problems during development may occasionally require the project to revert totally to an earlier phase. But overall, a step-wise process is generally followed.

In addition, this model should not be taken too literally. The labeling of phases does not mean that some parallelism is not possible and, in fact, common. Waiting for all testing to take place, for example, until the end creates grave risks to a project. Testing of various kinds may start in the earliest concept phase (perhaps in the form of simulations or prototyping) and usually continues through every phase. Also, various labels may be applied to the phases, depending on the type of system being designed. The goal is simply to define a structured process that will be tailored for the various needs of specific projects.

Why is all of this necessary for accident investigators and analysts to understand? While accidents occur during operations, the design flaws that led to them arise during development. The B737 MAX problems are an example of where it is not possible to simply blame the pilots (although people have tried). Accident investigations need to identify why the design flaws that led to the loss were introduced and why they were not identified before operational use of the system. Tracing the causes of accidents back to engineering and development is going to become more and more important as the complexity of the systems we are creating increases, making design flaws an important factor in accidents. To thoroughly understand why an accident occurred, we need to understand where any design flaws were introduced and why they were not identified before the system went into operational use. There are (or at least should be) many procedures designed into the development process that should find the design flaws before the system becomes operational. Investigating why such procedures were not effective is a large and important component of continual improvement in the development process.

Systems Theory vs. Complexity Theory

STAMP is based on systems theory, not complexity theory. Contributions to systems theory have been made by many people, but Ludwig von Bertalanffy, an Austrian biologist, is credited as one of the founders of what he called general systems theory. Norbert Weiner explored the implications for mathematics and engineering. Weiner called his theory *cybernetics,* but that term has faded over time and Bertalanffy's terminology is generally used today.

Systems theory, as you have seen in this handbook and perhaps elsewhere, is a set of principles that can be used to understand the behavior of complex systems, whether they be natural or man-made systems. *Systems thinking* is the term often used to describe what people are doing when they apply systems theory principles. These principles are briefly described in Chapter 1 of this handbook.

Complexity theory grew out of systems theory and other concepts in the 1960s and is usually associated with the Santa Fe Institute and researchers working there. Some commonalities exist between Systems Theory and Complexity Theory in that both include terms like emergence and focus on complex system behavior as a whole rather than on reduction or decomposition into components. The basic components of system theory are emergence, hierarchy, communication, and control, and these also are included in complexity theory. Both systems theory and complexity theory also include concepts

of feedback and feedforward control, adaptability, nonlinear interactions, and constraints. Both reject reductionism or decomposition as a principle for understanding system behavior.

There are, however, significant differences. Complexity theory was created to describe natural systems where seemingly independent agents spontaneously order and reorder themselves into a coherent system using laws of nature that we do not yet fully understand. Systems theory, in contrast, is more appropriate for man-made and designed systems where the system is purposely created by humans using some engineering or design process and where the basic design is known and changes are controlled.

Another major difference is that systems theory considers all systems as displaying emergent behavior while complexity theory divides systems into four types: simple, complicated, complex, and chaotic, each with different degrees or types of emergence.

As such, systems theory appears to be most appropriate for engineered or designed systems while complexity theory is most appropriate for natural systems where the design is unknown, such as the weather, and for sociological systems, such as communities, that are not designed and where there is a lack of order and it is very hard or impossible to predict the emergent behavior. Complexity theory deals with behavior that cannot be designed but instead must be experienced and studied as it arises.

Systems theory, at least for me, is easier to understand and use. Complexity theory frameworks can be difficult and confusing to apply and understand and may be overkill for the engineering goals related to improving safety and other emergent system properties. For this reason, systems theory was selected as the basis for STAMP.