Process Safety and Environmental Protection 212 (2026) 108886

Contents lists available at ScienceDirect

Process Safety and Environmental Protection

journal homepage: www.journals.elsevier.com/process-safety-and-environmental-protection

A powerful new approach to loss prevention☆

Nancy Leveson

MIT, 77 Massachusetts Ave., Cambridge, MA 02139, United States

ARTICLE INFO

Keywords:
Safety
Process safety
Systems theory
STAMP
STPA
CAST

ABSTRACT

Most safety engineering tools are 50–70 years old, but our technology continues to evolve, particularly with respect to the use of computers, the role of humans, and the scale of production. Tools, such as HAZOP, fault tree analysis, bow tie, and standard probabilistic risk assessment, were constructed using assumptions about the causes of accidents that are no longer complete. Successfully preventing losses in today's highly complex, tightly coupled, computer-intensive systems requires expanding these assumptions about causality and creating new safety engineering approaches based on them. This paper describes an expanded model of why accidents occur that is appropriate for today's systems. The new causality model treats safety as a *control* problem rather than simply a *failure* problem. New tools based on this new model allow modeling, analysis, and control of cultural and management as well as technical deficiencies. Scientific evaluations of new tools and comparisons with traditional tools are beginning.

1. Introduction: the problem

Technology has changed significantly since the time our current safety engineering tools were created. The traditional safety engineering tools are based on the assumption that accidents are caused by component failures (including human errors) or deviations in operating conditions that lead to component failures. They therefore treat the goal of safety engineering as preventing component failures.

This assumption and the efficacy of the traditional tools began to be challenged, starting in the 1970's, by the introduction of computer control, exponential increases in the complexity and coupling of our engineered systems and in the environment in which they operated, the introduction of new more powerful technology, and significant changes in the human roles in these new systems. Partly because of scale, managerial and organizations factors are assuming more importance in losses.

While accidents are still caused by component failures, new and different types of causes began to increase in importance. Analysis tools are needed that handle these new causal factors. To do so will require a paradigm change in how safety is handled in the process industry. This paper describes the changes that are necessary to prevent accidents in our more highly complex and tightly coupled process system designs. It also describes recent research comparing the new approach with HAZOP.

1.1. Limitations of traditional approaches

Safety traditionally has been viewed as a failure problem, where a chain of directly related failure events (and deviations leading to them) lead to a loss. To prevent accidents then, it makes sense to design to prevent individual component failures or to establish barriers between failure events in order to halt the chain of failures before a loss occurs.

To prevent or mitigate electromechanical device failures, safe design typically focuses on redundancy, overdesign, safety margins, etc. Defense-in-depth uses these design techniques to protect against multiple failures. To prevent human error, operators are provided with detailed procedures to follow with an emphasis on compliance.

Even in our older system designs, these approaches were not always successful. For example, common mode failures can defeat redundancy, as can other factors such as maintenance problems in critical equipment or design flaws in safety equipment. The Macondo accident involved problems in the design and maintenance in the blowout protector that defeated the use of redundancy (Baker, 2007). Bhopal had a large number of defense-in-depth measures that were undermined by inadequate maintenance procedures, flawed design, and economic pressures leading to some protection features, such as refrigeration, being turned off (Ayres and Rohatgi, 1987; Bogard, 1989).

As another example, the Texas City refinery explosion involved the use of a blowout drum that was not large enough and other outdated technology, inoperative alarms and level sensors in the ISOM process

☆ This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.
E-mail address: leveson@mit.edu.

https://doi.org/10.1016/j.psep.2026.108886
Received 30 December 2025; Received in revised form 14 April 2026; Accepted 16 April 2026
Available online 16 April 2026
0957-5820/© 2026 The Author. Published by Elsevier Ltd on behalf of Institution of Chemical Engineers. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

unit due to deferred maintenance and inadequate maintenance processes, and the inadequate design of safety devices (such as automated shutdowns or safety interlocks).

In addition to physical devices, safety at Texas City depended on correct and timely action by operators who were provided with inadequate procedures to follow, a control board design that did not provide accurate information for the operators, an emergency relief system that did not address the potential for a large liquid release in the event the raffinate splitter tower was overfilled, and trailers placed too close to a process unit handling highly hazardous material. Few of these causal factors involved random electromechanical device failures. Even those that did involve component failures were mostly due to inadequate maintenance procedures. These causal factors are not unique to the Texas City explosion but have been implicated in many recent process system losses (Baker, 2007).

Flawed management procedures and culture were also an important contributor to the events (Baker, 2007; Leveson, 2023). Cultural problems at Texas City reflected complacency about serious risks and low priority placed on process safety versus short-term profits. Management did not welcome or listen to reports of safety problems and ignored warning signs of changes that were leading to increasing risk. Company audits and assessments identified serious risks but did not lead to interventions, warning signs in the form of incidents were ignored, and serious deviations from safe operating procedures were tolerated. Cultural assessments at the Texas City refinery all agreed that personnel were not encouraged to report safety problems and were discouraged from reporting safety problems upward.

The organization and management of operations at Texas City contributed to the loss through diffusion of responsibility and authority and lack of corporate oversight and leadership.

While many losses today still result from component failures, new causal factors, such as interactions among system components that are all independently working as designed (as intended), are increasing in number and importance. Causal factors today include new types of human error, design errors, software requirements, cultural flaws, and poor management decision making (Leveson, 2012, 2023). Usually, all or many of these factors together contribute to the loss, changing the basic assumptions upon which hazard analysis techniques have been traditionally based.

Most of these new factors, or at least their new importance, are the result, starting in the 1970s, of changes in engineering. These changes include the introduction of computer control in nearly every system, exponential increases in complexity (partly allowed by the use of software), the introduction of powerful new technologies, and changes in the role of human operators.

The most significant factor is the introduction of new technology, particularly computer control. Computers do not "fail" in the same way as electromechanical system components. They do not wear out or randomly stop working, but they simply do something that was unsafe, where that is often recognized only in retrospect. The use of computers allows increased complexity, sometimes exponential, and less ability of designers and operators to understand and predict all the potential interactions among the system components. Increased coupling in our designs leads to lessened ability to predict all potential interactions in today's more complex and tightly coupled systems.

Human roles have also changed. Operators are less often directly controlling physical processes by following a set of procedures, but instead are engaging in the much more cognitively complex supervision and monitoring of automated controllers. They are not just making random errors, but are doing the wrong thing because they are confused about the current state of the system.

All of these changes are leading to a lowered ability to detect design errors before system use and to new causal factors in losses, changing the basic assumptions upon which traditional hazard analysis and design techniques are based. The events at Bhopal (Bogard, 1989), Texas City (Baker, 2007), Macondo (Graham and William, 2011), Fukushima Daiichi (Amano, 2014), and Shell Moerdijk (Dutch Safety Board, 2015) did not occur simply because of component failures, but instead they occurred because of a lack of effective *controls* on the physical system, management, and operating conditions. More powerful system analysis and design techniques, based on expanded assumptions about accident causality, particularly the role of inadequate control over the system behavior, are needed.

1.2. Levels of causation

Preventing losses today requires looking beyond technical deficiencies and to broadening our understanding of accident causes and our processes for preventing them.

In general, three levels of causes can be identified (Fig. 1).

The lowest level or "what" is comprised of the events and mechanisms that directly lead to the loss, for example, an explosion or an uncontrolled release of toxins. These events directly result from some "who" and "how" conditions (level 2), such as a bad valve design, software design flaws, or the operator not noticing something or not responding effectively to it.

The highest level is the systemic factors or "why." This highest level is comprised of the inadequate controls and design constraints that allow the occurrence of the Level 2 conditions. Examples include overconfidence and complacency; inadequate feedback, reporting, and learning from operations and incidents; inadequate management of change; production pressures; flaws in the system design process; poor maintenance procedures; and overreliance on redundancy or protection systems. A more complete list is provided in (Leveson, 2023) along with examples of how these factors have contributed to major accidents in the past.

The lowest level events and mechanisms are usually unique to a specific accident. The Level 2 conditions often apply to several accidents but within a limited context, while the Level 3 inadequate controls and constraints apply to a broad range and number of losses throughout an organization or industry and often involve indirect relationships with the level 2 conditions. Preventing the Level 3 or "systemic" causes will have the broadest impact on improving safety.

Loss prevention in the process industry, particularly the tools and design processes used, have historically focused on the lower two levels of causality. Making significant progress for the new types of systems and technology being developed and operated today, requires extending our modeling and analysis tools (1) to identify and prevent new types of level 1 and level 2 causes and (2) to identify and prevent level 3 systemic factors.

2. Extending the traditional approach to preventing accidents

The new approach to safety engineering does not throw out the old assumptions about the cause of losses but extends them. Rather than viewing safety as a *failure* problem, safety is viewed as a *control* problem. The goal is to control system behavior to prevent the system from getting into a hazardous state, that is, a state for which there are conditions under which a loss can occur. For example, in a chemical plant a system hazard might be a runaway reaction or an uncontrolled release of chemicals or toxins.

Control will still be needed over component failures and deviations in operating conditions, i.e., the traditional considerations, but control must also be provided over the other types of causal factors leading to hazards today. To accomplish this goal, simply improving our safety assurance and quantification techniques is not enough. Fundamental changes are needed, such as:

- More emphasis on up-front design of safety from the beginning of development.

  The primary emphasis in system development is too often on analyzing and assuring safety after much of the design and

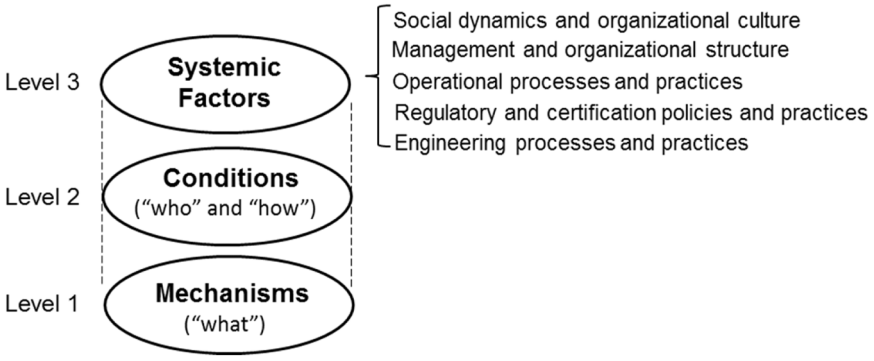


Fig. 1. Three Levels of Accident Causes (Leveson, 2023).

development has already been done. When systems were relatively simple and changes could easily be made, this approach was feasible. Emphasis could be placed on other system properties and then safety analyzed and improved if quantitative analysis showed the result was not acceptable.

The problem with this approach is that systems are now so complex that after-the-fact assurance is very difficult. It is not unusual for technical systems today, for example, to contain millions of lines of software. Not only is assurance of such complex systems infeasible, but making changes to the finished design is very likely to introduce other problems. It is almost always not practical to start over again or to make extensive changes.

The difficulty is exacerbated by the fact that the problems today are often in the requirements. Changing requirements late in development is extremely expensive and can cause the project to fail completely.

The resulting pressure *not* to find major safety problems is so high that often the goal of assurance is to show that the system is safe or the problems uncovered do not require extensive redesign. Safety assurance becomes a *pro forma* effort that leads to the acceptance of unsafe designs: Losses occur that could have been prevented. The crash of the British Royal Air Force Nimrod aircraft in Kandahar, Afghanistan in 2006 is unfortunately typical. One of the conclusions of the Nimrod accident report was that the focus in assurance was on showing the aircraft was safe rather then identifying potentially unsafe design features (Haddon-Cave, 2009).

Reliance on assurance is partly the result of the traditional hazard analysis techniques requiring a fairly detailed schematic or model of the system before they can be used. New approaches are needed that allow identifying and validating safety-related requirements at the beginning of a project and that support tradeoffs and decision making throughout design and development (concept of operations, requirement generation, architectural development), before a detailed system design has been created. In other words, these approaches would allow designing safety *into* the system from the beginning and providing assurance as design decisions are made.

- Better integration of human factors into hardware and software engineering.

The types of errors that humans are making today have changed. In the simpler, electromechanical systems of the past, human operators were directly controlling systems. They could be provided with pre-defined, step-by-step procedures to follow, with limited requirements for independent decision making.

In contrast, operators today are mostly controlling and monitoring the automation that is actually controlling the electromechanical components. Operator tasks require more cognitively complex problem solving that does not translate easily into a pre-defined procedure. This factor leads to requirements for *human-centered design* to reduce human errors and for more extensive knowledge and training of operators. The introduction of AI into process control systems will exacerbate all of these problems.

- More powerful hazard analysis techniques that include software and sophisticated human factors.

Attempts have been made to extend the traditional hazard analysis techniques to include new factors like software, but none have been successful in practice. Again, the problem is that the underlying causes of accidents are including more than just component failure.

Software does not "fail" unless the underlying computer hardware fails (which is rare today). Instead, software leads to losses when it does exactly what the developers intended it to do but that behavior is unsafe in the system in which the software is used. Indeed, most losses related to software today result from requirements errors (Leveson, 2012; Lutz, 1992): the role of software in losses results either from (1) incomplete or wrong assumptions about the operation of the controlled system or the required operation of the software or (2) unhandled controlled-system states and environmental conditions. We need new analysis approaches that include these new causes of losses. The design of the process control software used during the Shell Moerdijk explosion included features that had been purposely included but made the reaction harder to control by the human operators.

- New approaches to regulation and certification

Current assurance and certification processes assume that losses occur from component failures, namely, unreliability. The extensive use of software undermines these assumptions.

- Improved management of change (MoC)

Accidents usually follow some type of change, either in the system or in its operating environment. As our systems and the environment in which they operate increase in complexity, we will need improved management of change procedures, including unplanned changes.

For planned changes, our safety management systems need to incorporate stronger internal controls. Too often, defined management of change procedures are skipped or not implemented rigorously. In addition, the efficacy of the procedures used to analyze the safety of changes is reduced by the limited assumptions about the cause of losses.

For unplanned changes, more powerful leading indicators are needed. For example, leaks are commonly used as a leading indicator in the process industries. Leaks, however, would not have identified the increasing risk in the major process accidents of the past few decades. An approach to creating new *Assumption-Based Leading Indicators* has been proposed (Leveson, 2015). It is founded on the observation that most accidents result from changes in the assumptions— either planned or unplanned—upon which the safety of the design was based. Those assumptions are identified during the hazard analysis process, but the HA procedure used must include the ability to identify them.

3. STAMP: a new causality model for losses

As argued, many losses today stem from basing our attempts to prevent losses on a causality model that is too limited to fully explain losses today. Significant progress requires a paradigm change to a new, more powerful causality model. Such a model, called STAMP (System-Theoretic Accident Causality Model and Processes), serves as the foundation for new tools and approaches used in many industries today (Leveson, 2012). STAMP is based not on reliability theory but on Systems Theory.

The development of Systems Theory started about a hundred years ago in an attempt to create a science of complexity. A description of Systems Theory is beyond the scope of this paper, but it can be informally described as the assumption that "the whole is greater than the sum of its parts." The behavior of complex systems needs to be understood in terms not only of its individual components, but also the interactions among the components and the operation of the system as a whole. System properties, such as safety or risk, "emerge" from such interactions.

3.1. The STAMP causality model

To be successful, a new causality model needs to include the traditional chain of failure events model but go beyond it to include the causes of accidents resulting from our new technology and designs. In STAMP, losses are assumed to be the result of a complex, dynamic process. They arise in the interactions among humans, machines, and the environment. In essence, losses are treated as a dynamic control problem rather than a failure problem.

In STAMP, safety (preventing losses) requires enforcing a set of constraints on system behavior. For example, the level of pressure in the reactor must never exceed a particular value. Losses occur when the interactions among system components violate those constraints. Some of these constraints will include failures, such as valve failures. Others will involve other types of constraints, including constraints on design, management, and operations.

In STAMP, *risk* is defined not by the probability of failures and their potential severity, but by the power of the controls to enforce the safety constraints on system behavior. Risk is analyzed and evaluated using models of the system control structure, which consists of basic feedback control loops (Fig. 2).

A controller might be a human or automation. Controllers issue control actions based on their beliefs about the current state of the system, which is contained in a process model (a model of the current state of the process). When humans are the controllers, the process model is usually called a mental model. For software, the process model may be contained in software variables used by the software algorithm to make decisions about control actions. Today, the actuators and sensors may also be implemented by software. The process or mental model is updated by feedback from the controlled process.

Controllers use their process model to decide what control actions to issue to ensure that the behavior of the controlled process achieves the goals while satisfying the constraints. Note that this definition is just that for a standard feedback control loop. For example, when feedback indicates overpressurization in the controlled process, the controller may issue control actions to try to reduce that pressure. Decisions about how to respond are made by software algorithms or by various types of human decision making or a combination.

Software and human-related accidents can occur when the algorithm is incorrect or human decision making is flawed. It may also occur when the process or mental model is incorrect, that is, when the controller's internal model of the state of the controlled process is inconsistent with the real state. How can the latter happen? There are clearly many ways including the model being flawed from the beginning (e.g., poor training for humans or software flaws for automation) or feedback may be incorrect, missing, or delayed. Accidents can also occur, of course, when there are failures or design flaws in the controlled process or in the actuators or sensors.

Note that this model includes hardware, software, and human factors and acts as a common model on which all the system engineers can work and communicate. Companies have found that this feature helps solve many problems involving developer communication, particularly between human factors experts, software engineers, and traditional hardware and application experts.

Controllers include management, so the loop can be augmented with the basic organizational/management principles, shown in Fig. 3, of responsibility (attached to the controller), authority (represented by the assigned control actions), and accountability (implemented at least partly using feedback) when identifying hazardous management decision making and organizational design.

Of course, there is more than one loop in such a control structure. Fig. 4 shows an example hierarchical safety control structure divided into system development and system operation. Fig. 4 represents just an example of what functions can be included in the modeling and hazard analysis.

To assist in understanding a control view of safety versus a failure

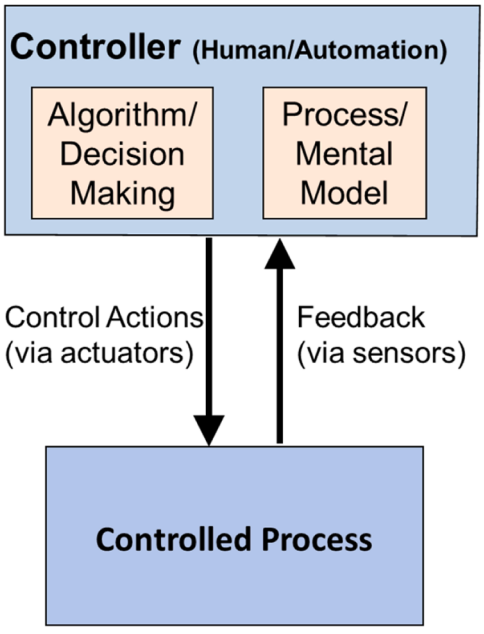


Fig. 2. The basic form from which control loops are created, consisting of a controller and a controlled process.

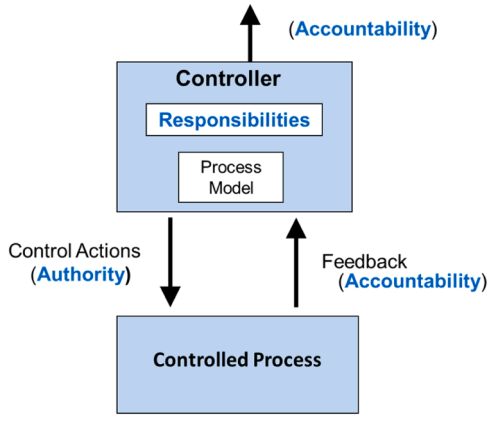


Fig. 3. A control loop augmented with the standard management controls of responsibilities, authority, and accountability.

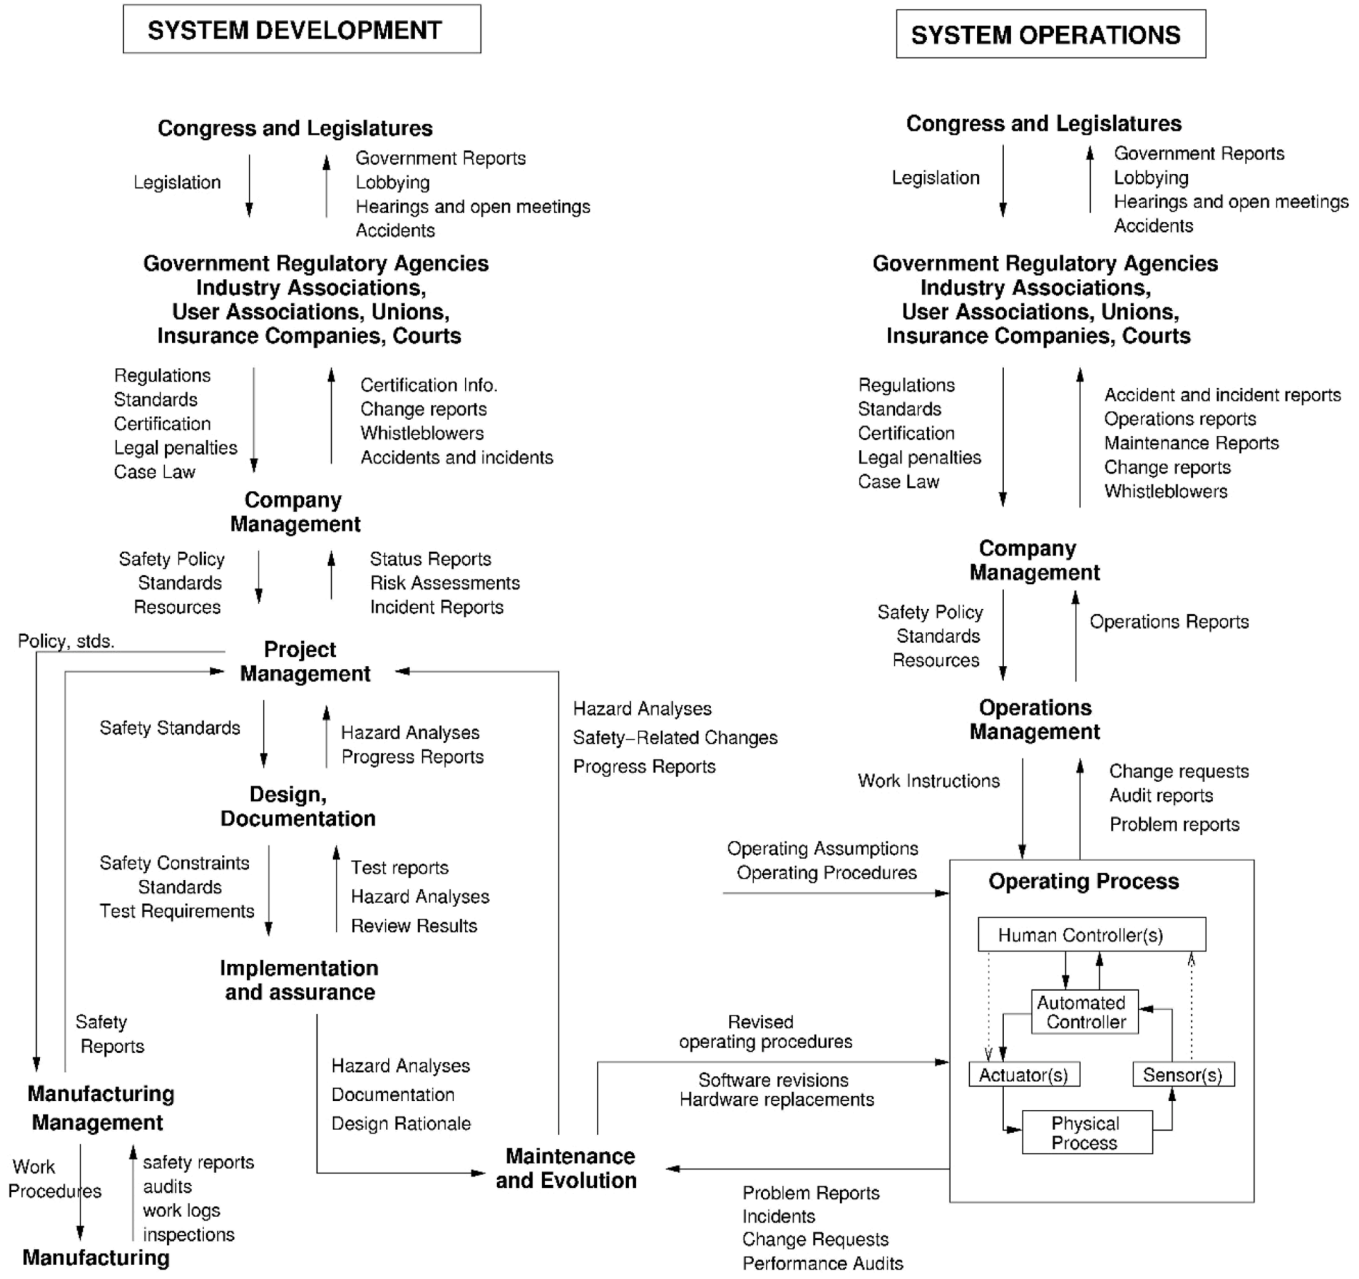


Fig. 4. Example of components that can be included in a hierarchical control structure.

view and how such a control structure can be used to investigate accidents or to design to prevent them, consider the Shell Moerdijk chemical plant explosion in the Netherlands in 2014 (Dutch Safety Board, 2015) and a causal analysis using STAMP (Leveson, 2017 and 2019). The original accident report was written from a failure point of view, however the Dutch Safety Board did an excellent job of going beyond simply listing the failures involved.

The Shell Moerdijk control structure is shown in Fig. 5. There are four general types of safety controls in this system: physical design controls, operational controls, management controls, and government controls (oversight).

Omitted from this figure (for simplicity) are the safety-related responsibilities of each component, available controls and feedback, and process models that need to be included.

Physical Controls: The Shell Moerdijk plant used standard defense-in-depth, i.e., a series of barriers to protect against runaway reactions, inadvertent release of toxic chemicals or an explosion (uncontrolled energy). Designed protections converted any released chemicals into a non-hazardous or less hazardous form, eliminated or reduced human or environmental exposure after release, and provided for the treatment of exposed individuals. In addition, the Shell Moerdijk plant had the standard types of safety equipment installed, such as alarms, a flare tower, pressure relief valves, etc.

The basic events (Level 1: "What") involved in the accident were the following:

1. The plant had been shut down for a short, scheduled maintenance stop (pit stop) to replace the catalyst pellets in the reactor and was being restarted at the time of the explosion.
2. One of the restart procedures is to warm up the reactors with ethylbenzene. During the warming (reheating) process, uncontrolled energy was released and unforeseen chemical reactions occurred between the ethylbenzene and the catalyst pellets that were used.

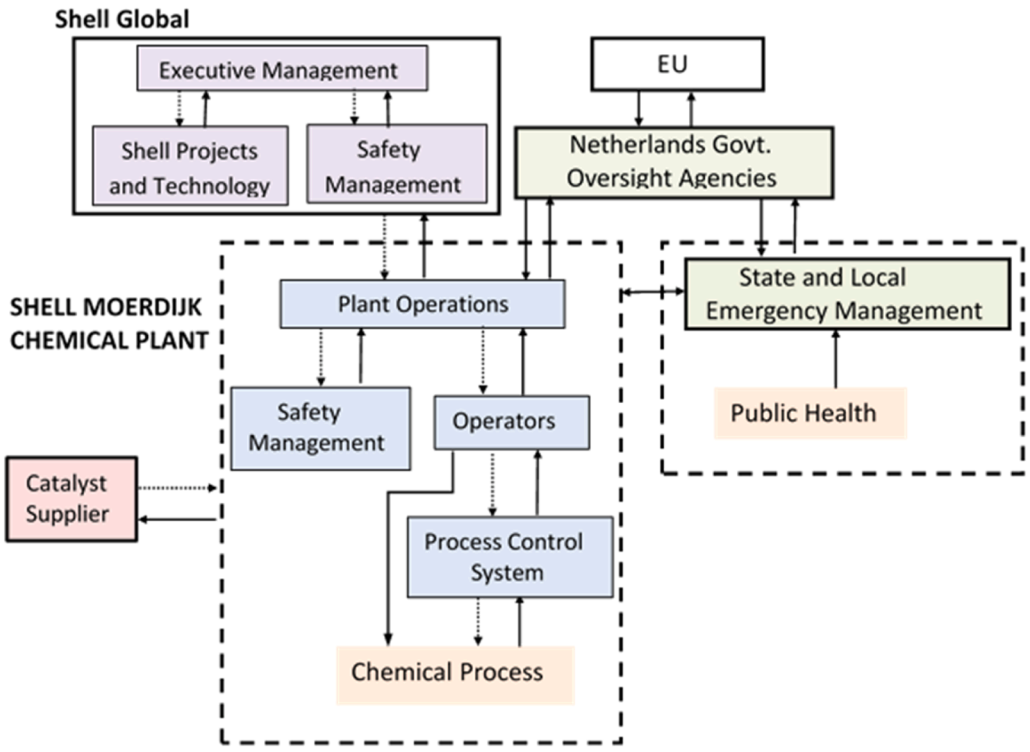


Fig. 5. The basic structure of the safety control structure for Shell Moerdijk.

3. The reactions caused gas formation and increased pressure in the reactors.
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.
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 an even higher increase in pressure.
6. The pressure rose so fast that it could no longer be controlled by the pressure relief devices, and the reactor exploded.
7. The contents of the reactor and its associated separation vessel were released into the wider environment.
8. Two people working opposite the unit involved at the time of the explosion were hit by the pressure wave of the explosion and the hot and burning pellets that were flying around.
9. A large raging local fire occurred, generating considerable amounts of smoke. Community firefighting, healthcare, crisis management, and crisis communications were initiated.

Note that no physical controls failed, such as the automatic protection system and the pressure relief valves: The safety features worked *as designed*, except for the final collapse of the reactor and separation vessel after pressure reached a critical level.

There were, however, several unsafe interactions among the system components. Many of these unsafe interactions were a result of design flaws in the reactor or in the safety-related controls. For example, due to unforeseen chemical reactions, hotspots were created that were not automatically detected due to an inadequate number of temperature sensors in the reactors. A runaway reaction caused gases to be released, increasing the pressure in the vessel.

At the same time, the maximum liquid level in a 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 flare (the exhaust gas system). This automatic protection device shutdown, however, had the unintended consequence of preventing the venting of the gas. The basic design flaw allowed the pressure to reach a point where the automatic pressure relief devices in place could not adequately release it.

A final example of a non-failure cause involves the placement of three remote-controlled containment valves that could not be remotely operated and thus had to be operated manually. Due to the intensity of the fire that had broken out and the risk of explosion, they could not be operated immediately.

Again, note that in this accident the protection devices worked as designed—they did not fail—but the design was flawed. Preventing future losses requires understanding more than just the causal factors related to the physical design.

Operational Controls: To demonstrate how the Shell Moerdijk safety control structure (see Fig. 5) can be used to understand why the automated and human controllers were ineffective in preventing this loss, consider Fig. 6 showing part of the operational control structure and causal analysis.

The automated process control system was configured for normal operation and provided only minimal to no assistance during startup, when the operators most needed it. It did not stop the process when pressure and temperature increased to unsafe levels and did not provide an automated reset after two high-high level alarms caused the automated protection system to close the gas discharge system. Note, however, that the reactor provided incomplete, missing, or erroneous temperature and was not capable of detecting the hotspots that occurred. This bad temperature and pressure information was provided to the operators. At the same time, some important information was not provided to the operators, such as the rate of the temperature increase.

The operator actions, in hindsight, appear to give them major responsibility for this loss. They did not properly adjust gas and liquid flows, increased heat when it was already too high, followed the incorrect and incomplete work instructions they were given, and did not activate the manual trip switches or respond to alarms. Most accident reports label these as software and human “failures” and stop there.

Most of the operators’ actions can be understood, however, by

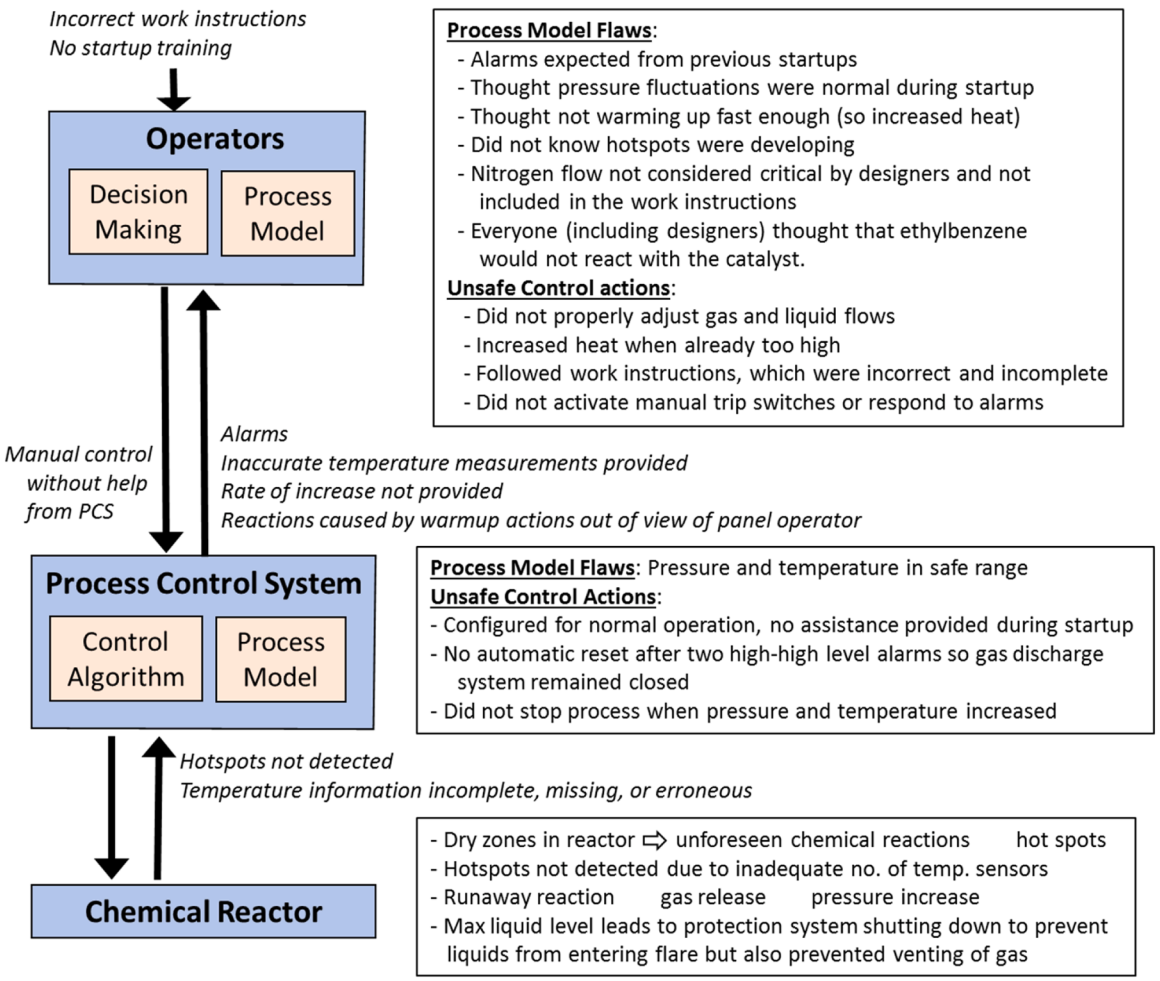


Fig. 6. Partial Explanation of Shell Moerdijk Plant Explosion.

looking at their process models and the context in which they were making decisions. For example, they were not trained in starting up the plant and did not have assistance from the automated process control system. The alarms and pressure fluctuations were expected and thought to be normal. They also thought the reactor was not starting up fast enough: they were not provided with information about the rate of increase, which is difficult for humans to determine under the conditions in which they were working so they increased heat. Basic flawed assumptions about the process were common to both the operators and the designers, for example that nitrogen flow was not critical (so not included in the work instructions). Everyone believed that ethylbenzene would not react with the catalyst and therefore that a runaway reaction was impossible.

Although control of startup for the Moerdijk plant was difficult, a decision was made not configure the process control system to assist the operators or to stop the process when the pressure and temperature increased precipitously or to issue an automatic reset after two high-level alarms, so the gas discharge system remained closed. No reason is given in the accident report for the decision for the automated process control system not to provide operator assistance during the restart phase of operations or why an emergency stop capability was not designed into the system.

Only one small part of the causal analysis is shown in Fig. 6, but the entire control structure has been analyzed using STAMP-based tools (Leveson, 2017). The analysis of the upper parts of the control structure identify the systemic (Level 3) factors involved in management and governmental controls. For practical reasons, only parts can be included here.

Safety Management Controls: Some of the reasons for the inadequate safety management controls can only surmised from the accident report. If a control view of the accident had been used in the original investigation, human designers and controllers could have been asked why they made decisions about the plant safety controls that contributed to the loss. A very large number of flawed decisions contributed to the loss due to inadequate management controls. A few examples are included here.

The designers of the plant did not recognize or understand the risks involved in the reactor startup process so those 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.

Why did the design flaws get through the plant design and review process? Little information was provided in the accident report about that review process. It is an important control that needs to be considered in understanding why a loss occurred.

Some of the design flaws resulted from changes over time, such as changes in the composition of the catalyst that originally had been scientifically proven to be nonreactive with ethylbenzene. Although the composition of the catalyst changed over time, no further testing was ever done. Control over this potential risk could have been could have been provided by the adoption of a leading indicator of increasing risk that checked for changes in the catalyst. Such leading indicators can be created from the original hazard analysis (Leveson, 2015).

In general, the Shell management of change procedures were not followed or used effectively. Risks resulting from changes made to the

plant, the catalyst, the processes, and the procedures were not identified and controlled. In fact, the change in the catalyst had led to similar incidents at other Shell installations but no explosion resulted because of special conditions existing at those instances.

The incorrect mental model of the Moerdijk hazard analysts was an important causal factor in this accident. Was the information that the catalyst was no longer nonreactive with ethylbenzene not included in the Shell safety information system after the prior incidents? How is such critical information controlled in the company? Was the information available but the Moerdijk designers never checked it? Did they know about it but not think it was applicable to their plant? Understanding why standard controls have not been effective is important in reducing similar hazards in the future.

As just one more example, operations management at Shell Moerdijk did not provide appropriate training for the operators for both nominal and off-nominal control activities. The work instructions for the reactor startup were written by the operators, who did not have the experience required to write such instructions, and were never reviewed by safety engineering experts. In addition, operations management did not assign operators to the startup process that had the qualifications stated in the refinery safety report.

Governmental Controls/Oversight:

Dutch regulatory agencies provided oversight and control over hazardous facilities including the Shell Moerdijk chemical plant. At least partly because of limited resources, the government used "performance-based regulation" that placed responsibility on the operators of high-risk activities to identify their own shortcomings. During annual inspections, regulators only ensured that companies had the right documented procedures along with a few spot checks to determine whether they were being used. During inspections, regulators gave Shell Moerdijk a pass on behavior that might have instead been labeled as violations. They did not notice or did not react to Shell not acting in accordance with its own safety management system.

3.2. *Using STAMP*

STAMP is a theoretical causality model, not a tool. Tools are needed to employ the theoretical causality model in engineering practices. Fig. 7 shows the tools and processes that have been created using STAMP as a theoretical foundation.

Hazard Analysis: STPA (System-Theoretic Hazard Analysis) (Leveson and Thomas, 2018) is based on STAMP. STPA, like the traditional hazard analysis techniques, is a step-by-step process to identify how losses might occur so that they can be prevented. Fig. 8 shows the four steps.

First, the losses and related hazards are identified and categorized by system stakeholders. Then the system safety control structure model is created. The analysis is performed on this control structure, first identifying potential unsafe control actions that could lead to hazards. Then the scenarios that could lead to these control actions are identified. The causal scenarios include not only the failure and deviation scenarios found by the traditional hazard analysis techniques, but also design flaws, software flaws, sophisticated human factors problems, flawed management decision making, etc.

STPA is widely used in the aviation, defense, and automotive industries and international standards have been created. It can be used in the early phases of system engineering, such as in the generation of a CONOPS (Concept of Operations) and system requirements, and used to create safer system designs and operational practices from the beginning of development. It can also be used after a detailed design has been created, but the options to eliminate or mitigate unsafe behavior will be more limited at that time.

Accident Analysis: STAMP provides the basis for an accident analysis technique called CAST (Causal Analysis based on Systems Theory). CAST provides more information about causality, particularly the systemic factors, than most accident reports provide today. CAST was used for the causal analysis of the Shell Moerdijk explosion (Leveson, 2017) partially described in this paper and for many other losses in a variety of industries.

Security Analysis: STAMP is applicable to security as well as safety. STPA can be used to identify loss scenarios that involve intentional causes of loss as well as unintentional ones. STPA does not distinguish between them, it simply identifies the scenarios that can lead to a "loss" so that measures can be taken to eliminate them through the system design process and through operations practices.

Identifying Leading Indicators of Increasing Risk: While lagging indicators, such as leaks, are used in the process industry, identifying leading indicators has proven to be more difficult. Risk increases after changes occur that violate the original design and analysis assumptions. The STPA analysis embodies those assumptions and can be used to identify leading indicators (Leveson, 2015).

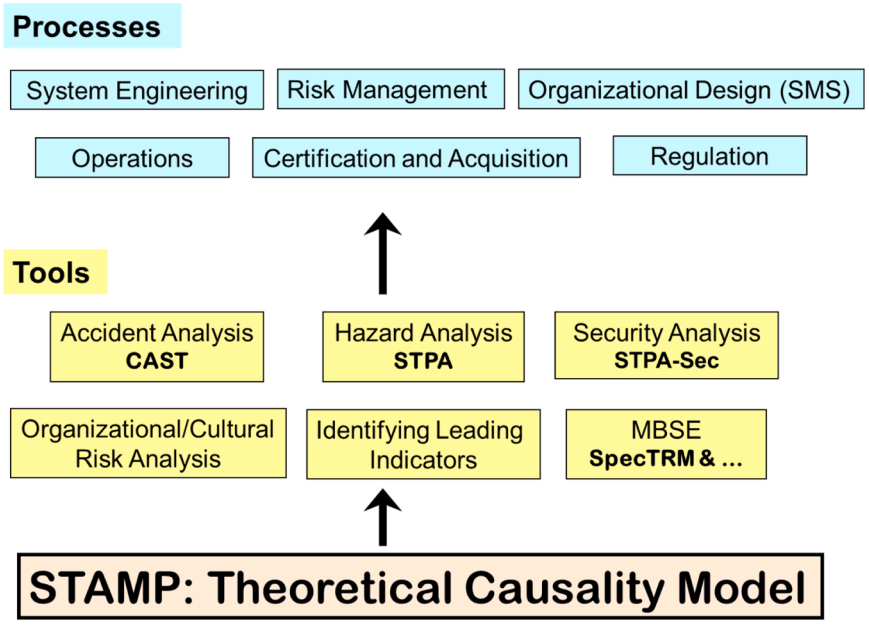


Fig. 7. The tools and processes that have been built on the theoretical STAMP causality model.

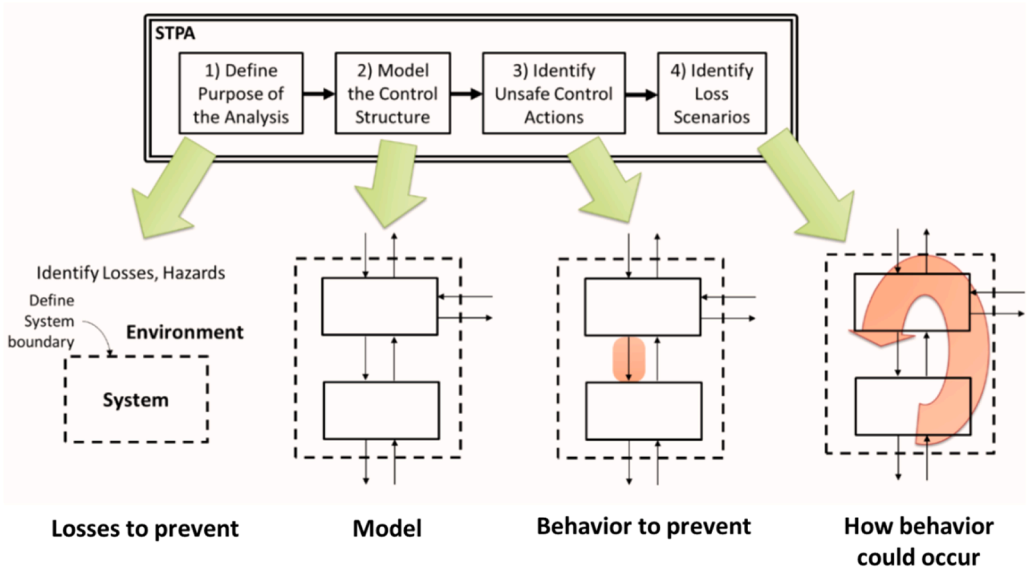


Fig. 8. The steps in performing STPA.

3.3. Other uses

STAMP (and System Theory) provides a general causality model that can identify potential causes of many types of losses, not only those usually associated with safety. For example, most any type of loss can be analyzed such as environmental pollution, property damage, mission loss, or negative business impact (such as damage to reputation). STPA has also been used to improve organizational design, certification and regulation, risk management, operations, and the general system engineering process.

4. Evaluation and comparison with HAZOP

STPA is widely used in aviation, automotive, and defense, but less so in the process industry, where the use of HAZOP is common. HAZOP is a very powerful hazard analysis method that looks not just at failures (such as FTA and Bow Tie) but at deviations. It is informative to examine the relationship between STPA and HAZOP. Very few scientific comparisons have been made. A common problem with published comparisons is that while the HAZOP teams are usually expert in using the technique, the STPA teams, because the technique is relatively new, often are inexperienced and the resulting STPA analysis is not performed well. In addition, the applications used are often very simple. The published comparisons also have other methodological weaknesses.

The most scientific and useful comparison with HAZOP involves a process industry application involving a fertilizer plant (Jamal et al., 2022). STPA and HAZOP were used by independent and qualified expert teams to discover flaws in a real system. Neither team had any preexisting knowledge of flaws before applying the methods. The system contained real design flaws that had led to the incident during the operation of the system. The outcomes and recommendations of HAZOP and STPA were compared to determine what differences exist, if any, and identify whether critical gaps exist for modern process industry applications. The HAZOP and STPA results are also compared to the corrective actions produced after a hazardous and costly incident during operation.

A real fertilizer plant application with both automated and human controls was selected for the comparison. The application was analyzed by a qualified HAZOP team before commissioning and operation. Later an incident occurred that had not been identified by the HAZOP performed before operation of the plant. STPA was performed on the same system to identify the weaknesses. Each team included experts with appropriate training and over a decade of industry experience applying the selected analysis method.

The application contained weaknesses and flaws that were recognized after it was put into operation. Both teams were unaware of the weaknesses and operational events and had no knowledge of either when they performed the hazard analyses. Both teams were limited to general design information available before the application was put into operation.

The system analyzed was a CO_2 removal system with automated and human controls. The HAZOP analysis was performed following industry standard recommendations and was led by HAZOP facilitators with decades of expert professional HAZOP experience. The HAZOP team was a multi-disciplinary team that was experienced and fully qualified to perform HAZOP. All HAZOP results were independently reviewed to identify possible gaps or omissions.

The HAZOP produced useful insights about potential hazards that had to be mitigated. A total of 436 actions were generated, with the majority related to interlocks, alarms, control actions and procedural control, “no flow,” and many other conditions, as per standard practice. The team used a rigorous methodology, including additional LOPA-style methodology. However, a comparison of the HAZOP results and the actual operating experience, which was not available to the HAZOP team, showed that the HAZOP did not identify the mission requirement of seal gas supply valve cutoff during Methanator tripping that contributed to the hydrogen fire.

STPA was applied on the same system and with qualified STPA practitioners. As with those who performed the HAZOP, the STPA users did not know anything about the incident that had happened, the root causes, or the necessary corrective actions to address the hidden weaknesses in the system.

The researchers made the following conclusions about the comparison: Both HAZOP and STPA produced useful insights about potential causes of hazards and losses. Common results related to physical behaviors, failures, and deviations in the controlled processes were identified by both methods, such as a valve blockage or failure. The most significant difference is that the HAZOP results—performed by expert practitioners and with comprehensive reviews—did not adequately consider:

- Weaknesses in the intended design of the automation;
- Missing logic or unsafe logic in automated controllers;
- Safety-critical functionality missing from the design;

- Missing feedback information to enable safe decision-making by controllers;
- Interactions between multiple controllers;
- Incorrect operational assumptions made by human operators, especially assumptions about the automation;
- Inconsistent automation that will cause human operator confusion, like automatic mode changes or an automatic trip function that works in almost every possible case except one;
- Gaps in human procedures or training;
- Assumptions about human operator behaviors that were not validated;
- Potentially conflicting commands from multiple controllers, such as a condition that simultaneously triggers an automated downstream trip and a manual human command to open a valve; and
- Common causes that defeat assumptions about redundancy and independence.

The STPA method was found to capture hazardous human and automation-related behaviors that were missed by HAZOP, and STPA generated critical recommendations missed by HAZOP that would have prevented the real adverse events. Only the STPA results anticipated the causes and corrective actions that were discovered after the hazardous and costly event during system operation.

The researchers note some differences in the teams used in the study. The STPA team was less than half the size of the HAZOP team. In addition, the time allocated to perform STPA was less than half that of HAZOP. The HAZOP team consisted of many domain experts who were fluent in the detailed design and operation of the system, while the STPA team had less familiarity with the system. The STPA team did have the ability to ask specific questions about the system if the questions were generated by the analysis. All information and answers were limited to materials available before the operation of the system and the incidents that occurred.

5. Summary

HAZOP is a very powerful hazard analysis technique that has served the process industry well for 60 years. However, engineering is changing. Complexity in the systems we are designing and operating is reaching a new level, i.e., a tipping point, where the old safety approaches are becoming less effective. New causes of losses are appearing.

Traditional hazard analysis and accident analysis techniques developed in the 1960s do not provide all the information necessary to prevent the new causes of losses in these systems. In the fertilizer comparison study described in Section 4, the study authors concluded that HAZOP provided useful insights, but even when performed by expert practitioners with multiple reviews and in compliance with industry standards, HAZOP was found to overlook important causes related to new technology, human behavior, and interactions between non-failed components operated as designed. STPA addressed the gaps observed in the HAZOP method, including identifying real scenarios and real corrective actions for costly losses that actually occurred during operation of the plant.

STAMP represents a paradigm change from current reliability-based approaches to a “systems approach” based on Systems Theory. It changes the focus from increasing component reliability—i.e., preventing failures—to enforcing constraints on system behavior through various types of controls. Such a change allows creating new analysis and engineering design approaches that are more powerful and inclusive, have been found to be less expensive, and that work on extremely complex systems. This new approach can be used to design safety, security, and other system properties into the system from the beginning of development, even before a detailed design exists.

Experience so far, mostly in the aviation, defense, and automotive industries, shows this new approach is much more effective in today’s complex, software-intensive systems. Scientific comparisons and evaluation in the process industry are only beginning. Additional scientific comparisons are needed, including the ability of trained analysts to perform the analyses and relative costs. Comparisons of costs in other industries have found STPA to be less expensive than other standard techniques, particularly after training and experience in using STPA. But comparisons to HAZOP are needed. Additions to STPA that include characteristics unique to the process industry may be useful.

Even if the answer is not STAMP and its tools, the process industry needs to figure out how to upgrade the hazard analysis tools and processes it uses to deal with the changes in technology and new causes of losses that are arising over time.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Amano, Yukiya, 2014–2025. The Fukushima Daiichi Accident: Report by the Director General. IAEA (International Atomic Energy Agency). ⟨https://www-pub.iaea.org/mtcd/publications/pdf/pub1710-reportbythedg-web.pdf⟩ (available for download from).

Ayres, Robert U., Rohatgi, Pradeep K., 1987. Bhopal: Lessons for technological decision makers. In: Technology in Society, 9. Pergamon Journals, pp. 19–45.

Baker, James (chair) (2007) The Report of the BP U.S. Refineries Independent Safety Review Panel, January, download available from ⟨http://news.bbc.co.uk/2/shared/bsp/hi/pdfs/16_01_07_bp_baker_report.pdf⟩ (last accessed 12/29/2025).

Bogard, William, 1989. *The Bhopal* Tragedy. Westview Press, Boulder CO.

Graham, Bob and William K.Reilly (co-chairs) (2011) Deep Water: The Gulf Oil Disaster and the Future of Offshore Drilling, Report to the President: National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling, January, available for download from ⟨https://www.govinfo.gov/content/pkg/GPO-OILCOMMISSION/pdf/GPO-OILCOMMISSION.pdf⟩ (last accessed 12/29/2025).

Dutch Safety Board, 2015. Explosions MSPO2. Shell, Moerdijk, The Hague.

Haddon-Cave, Charles, 2009. The Nimrod Review, 1025. The Stationery Office Limited, London (HC).

Jamal, Kamran, Faisal, Arooba, Arif, John, Arooj, Thomas, P., 2022. Comparative Analysis of Hazard and Operability Study (HAZOP) and Systems Theoretic Process Analysis (STPA). 5E AIChE.

Leveson, Nancy, 2012. Engineering a Safer World, 2012. MIT Press, Cambridge, MA.

Leveson, Nancy, 2023. An Introduction to System Safety Engineering. MIT Press, Cambridge, MA.

Leveson, Nancy and Thomas, John (2018), *STPA Handbook*, available for download from ⟨http://psas.scripts.mit.edu/home/books-and-handbooks⟩ (last accessed 12/29/2025).

Leveson, Nancy G., 2015. A systems approach to risk management through leading indicators. Reliab. Eng. Syst. Saf. 136, 17–34. https://doi.org/10.1016/j.ress.2014.10.008.

Leveson, Nancy G. (2017) *CAST Analysis of the Shell Moerdijk Accident*, available for download from ⟨http://sunnyday.mit.edu/shell-moerdijk-cast.pdf⟩ (last accessed 12/29/2025).

Leveson, Nancy G. (2019) *CAST Handbook*, available for download from ⟨http://psas.scripts.mit.edu/home/books-and-handbooks⟩ (last accessed 12/29/2025).

Lutz, Robyn R., 1992. Analyzing software requirements errors in safety-critical embedded systems. Proceedings of the International Conference on Software Requirements. IEEE.