A More Powerful Approach to Process Safety

Nancy G. Leveson
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology

Accidents in the process industry continue to occur and great progress is not being made in reducing them. Post mortem analysis usually indicates that they were preventable and had similar systemic causes. Why do we fail to learn from the past and make adequate changes to prevent reoccurrence? Why are we not successful in the first place?

A variety of causes for why accidents occur have been offered: human operator errors, component failures, lax supervision of operations, poor maintenance, etc. All of these explanations, and many others, have been exhaustively studied, analyzed, “systematized” into causal groups, and a variety of approaches have been developed to address them. And yet they still occur with significant numbers of fatalities and injuries, with disruption of productive operations, and frequently with extensive destruction of the surrounding environment, both physical and social.

Is it true that the problem of ensuring safe process operations is so complex that it defies our technical and managerial abilities to address it successfully? Should we consider process accidents as "normal" events and factor into our considerations the cost of addressing their consequences (as is common today)? Or is it that we are not going about trying to prevent them in the most cost-effective way?

Process systems today are changing in their nature. They are more complex, partly due to the introduction of computers, particularly to monitor and control physical processes. To prevent accidents in these increasingly complex systems, a new approach is necessary. The new “systems approach” redefines the safety problem as not just preventing system component failures but instead imposing constraints on the behavior of the system as a whole to prevent accidents.

The systems approach has resulted in much greater ability to design systems to prevent accidents and to learn more from the accidents and incidents that do occur. It is not new nor untested, although it is new to process safety. It is now the primary approach to the design of automated vehicles and is being used on some of the most complex systems being built today in the fields of aviation and defense.

Surprisingly, this new approach, although much more powerful, is turning out to require fewer resources than our current, less successful techniques. Return on investment (ROI) information is just beginning to be available, but companies are reporting large savings. It can be adapted to be used in process safety.

Accidents in the process industry continue to occur, and near misses are multiplying in alarming numbers, for two basic reasons: (1) The traditional analysis methods used do not discover all the underlying causes of events and (2) Learning from experience is not working as it should. A common reason for both of these is rapid changes in engineering practice that is making our old approaches less effective. To make progress, we need to re-examine the entrenched beliefs, assumptions and paradigms that underlie process safety engineering to identify disconnects with the prevailing methods.

This short paper offers a system theoretic view and approach to a rational, all-encompassing new framework for addressing process safety. It is based on a control-inspired view of the process safety problem, which is amenable to modern model-predictive control approaches, and can encompass all potential causal factors in accidents—from those on engineered systems at the processing level to those associated with operations management, regulations by governmental agencies, standards by insurance companies, and legislation governing operating plants.

The white paper starts by questioning the applicability the traditional accident causation models, which have constituted the basis for the development of almost all process engineering tools dealing with process safety. These include HAZOP, fault-tree analysis, FMEA, bow-tie analysis, etc. It then describes an alternative "violation of safety constraints" as the fundamental underpinning of a comprehensive framework for process safety engineering, both during development of a new process and during operations. Process safety is redefined as a *system problem* by underlining its fundamental distinction from reliability engineering.

1. The Prevailing Accident Causation Models and Their Shortcomings

In process safety the prevailing assumption is:

Accidents are caused by chains of directly related failure events.

This assumption implies that working backward from the loss event and identifying directly related predecessor events (usually failures of process components, or human errors but also deviations of process variables) will identify the "root cause"[1] for the loss. Using this information, either the "root cause" event is eliminated or an attempt is made to stop the propagation of events by adding barriers between events, by preventing individual failure (or deviation) events in the chain, or by redesigning the system so that multiple failures are required before propagation can occur (putting "*and*" gates into the event chain). Figure 1 shows a typical chain-of-events model for a tank rupture. The events are annotated in the figure with standard engineering "fixes" to eliminate the event.

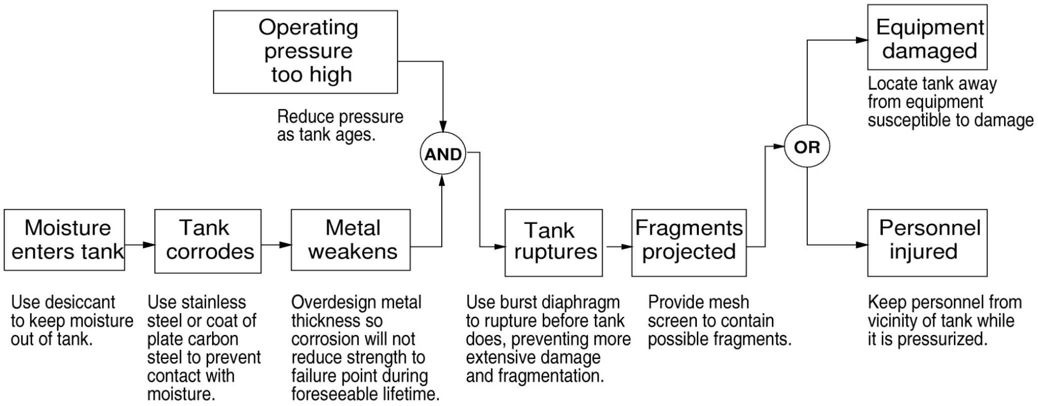

Operating pressure too high
Reduce pressure as tank ages.
AND
Equipment damaged
Locate tank away from equipment susceptible to damage
OR
Moisture enters tank
Tank corrodes
Metal weakens
Tank ruptures
Fragments projected
Personnel injured
Use desiccant to keep moisture out of tank.
Use stainless steel or coat of plate carbon steel to prevent contact with moisture.
Overdesign metal thickness so corrosion will not reduce strength to failure point during foreseeable lifetime.
Use burst diaphragm to rupture before tank does, preventing more extensive damage and fragmentation.
Provide mesh screen to contain possible fragments.
Keep personnel from vicinity of tank while it is pressurized.

Figure 1: An Example Chain-of-Events Model for a Tank Rupture

In this example, the chain of events is identified as moisture entering a tank leading to tank corrosion leading to weakened metal. The weakened metal, along with a specific operating pressure, leads to the tank rupturing and then fragments being projected. The final loss would be personnel injuries and/or equipment damage. To reduce or eliminate accidents, each event could either be made less likely to occur or the propagation of failures could be stopped. The annotations on the events show a typical design or operational solution to preventing the final loss.

Such linear, chain-of-event accident causation models constitute the foundation on which the process safety methodologies and engineering tools, such as Bow Ties, Hazard and Operability Analysis (HAZOP), Fault Trees, and Failure Modes and Effects Analysis (FMEA). However, this model has several serious drawbacks, including the oversimplification of causality and the accident process, the exclusion

[1] Note that the identified root cause is arbitrary, depending on how far one follows the causal chain or what events are considered and included.

of many of the most important factors in accidents, and incomplete consideration of the role of the social components of systems and not just the technical ones.

Oversimplification of causality and the accident process:

Most current accident models and accident analysis techniques suffer from the limitation of considering only the events underlying an accident and not the entire *accident process*. In particular, they omit *why* the events occurred beyond direct causality with the immediately previous event.

For example, why might the tank corrode even if desiccant is used? One reason might be that operating personnel forget to add desiccant or a new tank is installed using a different type of metal. Why might the operating pressure not be reduced as the tank ages? Perhaps the company is in a financial and competitive situation where productivity cannot be reduced at this time and reduction in pressure is put off until a future time. Why might operating or maintenance personnel find a need to enter the screened off area around the tank without shutting down operations completely during that period of time? There are lots of potential reasons including the cost of shutting down production for what appear to be quick and simple chores.

These additional causal factors are not events: they explain why the events occurred despite the precautions (as shown in the annotations) that have been taken to prevent them. The potentially most effective solutions are not suggested by the chain of events because the potential causal factors are more complex than suggested by the simple event chain model.

Also, within the scope of the prevailing chain-of-events approach, it is usually very difficult if not impossible to find an "event" that precedes and is causal to observed operator behavior. For example, it is nearly impossible to find a clear link between the design features of the processing units or automated controllers and operator actions. Furthermore, instructions and written procedures on how to start-up a plant or switch its operation to a new state are almost never followed exactly, as operators strive to become more efficient and productive, and deal with time constraints and other pressures.

Humans do not just make random errors. Their behavior is influenced by the design of the system in which they are working. Humans are limited by the physical controls provided, by the type and amount of information they have about the state of the process being controlled, and by management pressures and the culture of the company in which they work.

Finally, systems are not static, including process systems. Rather than accidents being a chance occurrence of multiple independent events, they often tend to involve a *migration to a state of increasing risk over time* as human behavior and the environment changes. A point is reached where an accident is inevitable, unless the new higher risk is detected and reduced. The particular events involved are somewhat irrelevant: if those events had not occurred, something else would have led to the loss. This concept is reflected in the common observation that a loss was "*an accident waiting to happen.*" Behavior will change over time, perhaps as the result of a search for greater efficiency, profits, or the competitive nature of the business environment at the time.

Understanding and preventing or detecting system migration to states of higher risk requires that our accident models consider the *processes* involved in accidents and not simply the events and conditions: Processes control a sequence of events and describe system and human behavior as it changes and adapts over time—perhaps as a result of feedback or a changing environment—rather than simply considering individual physical events and human actions.

Exclusion of many of the systemic factors in accidents, indirect or non-linear interactions among events, software, and human factors:

Accident causation models oversimplify causality because they exclude systemic factors—such as budget cuts due to poor business conditions, the desire to increase productivity, conflicting goals, or difficulty in finding appropriate parts or well trained personnel—which have only an indirect relationship to the events and conditions in the chain-of-proximate events. A few attempts have been made to include systemic factors in such models, but they are severely limited in achieving this goal.

Not accounting for the role of management, regulation, and nontechnical factors:

Accident causation is a complex process involving the entire sociotechnical system including legislators, government regulatory agencies, industry associations and insurance companies, company management, technical and engineering personnel, operators, etc. To understand why an accident has occurred, the entire process needs to be examined, not just the proximate events in the event chain. Otherwise, only symptoms will be identified and fixed, and accidents will continue to recur.

Jerome Lederer, the "father of aviation safety," observed that system safety should include non-technical aspects of paramount significance in assessing the risks in a processing system:

> "System safety covers the entire spectrum of risk management. It goes beyond the hardware and associated procedures to system 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 attitudes 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."

Belief that accidents are caused only or primarily by component failures:

An underlying assumption today is that

Process Safety is increased by increasing the reliability of the individual system components.

This assumption essentially concludes that if components do not fail, then accidents will not occur.

However, a high-reliability chemical process, i.e. a process with highly reliable engineered components (vessels, piping, connectors, pumps, sensors, control valves, control algorithms, etc.) is not necessarily a safer process. Unsafe interactions among the process components can lead to unsafe operations, while all components are functioning as intended.

As an example, consider the case of an accident that occurred in a batch chemical reactor in England (Kletz, 1982). The design of this system is shown in Figure 2. The computer was responsible for controlling the flow of catalyst into the reactor and also the flow of water into the reflux condenser to remove heat from the reactor. Sensor inputs to the computer were provided as warning signals of any problems in various parts of the plant. The specifications of the monitoring and control algorithms required that if a fault was detected in the plant, the controlled inputs should be left at their current values, and the system should sound an alarm.

On one occasion, the computer received a signal indicating a low oil level in a gearbox. The computer reacted as its specifications required: It sounded an alarm and left the controls as they were. This occurred when catalyst had just been added to the reactor and the computer-based control logic had just started to increase the cooling water flow to the reflux condenser. When the computer stopped and left the controlled components at their current values, the cooling water flow was kept at a low rate. The reactor overheated, the relief valve lifted, and the contents of the reactor were discharged into the atmosphere.

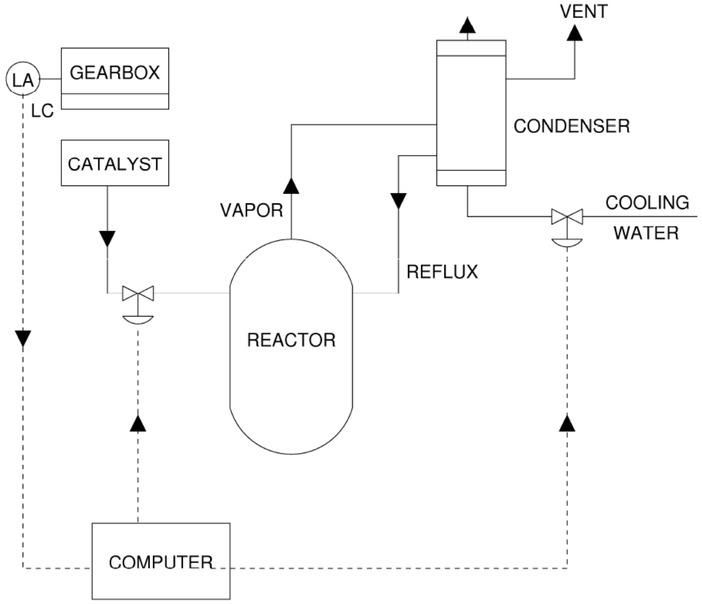

Figure 2. Batch reactor system

There were no component failures involved in this accident: the individual components, including the software, worked as specified and intended by the system designers, but together they created a hazardous system state. Merely increasing the reliability of the individual components or protecting against their failure would not have prevented the loss. Prevention required identifying and eliminating or mitigating unsafe interactions among the system components.

Indeed, most software-related accidents have been this type of system design accident—they stem from the operation of the software, not from its *lack* of operation and usually that operation is exactly what the software engineers intended. Thus chain of failure events models as well as system design and analysis methods that focus on classic types of failure events will not apply to software-intensive systems.

As the operation of petrochemical processes is increasingly controlled by software systems, the ensuing complexity in identifying and eliminating potential accident causal factors is increasing exponentially. This phenomenon is what I have called the *curse of flexibility*. Physical constraints restrict the values of physical quantities and thus impose discipline on the development, design, construction, and operation of a chemical plant. They also control the complexity of the processing plants that are being built. With control software, we can simultaneously vary the values of hundreds of flow-controlling valves to regulate the values of hundreds of measured outputs, and with modern real-time optimization algorithms we can optimize process operations by varying the values of hundreds of control set points.

What is possible to accomplish with software systems, however, is different than what can be accomplished *successfully* and *safely*. The limiting factors change from the structural integrity and physical constraints on materials to limits in human intellectual capabilities. The accident caused by the correct deployment of a software-based control system in the batch reactor of Figure 2 is a manifestation of the dangers lurking in the increasing usage of software systems in the control and optimization of processing operations. In this example, the primary concern is not the "failure" of the software control system, but the lack of appropriate safety constraints on the behavior of the software system. Clearly, the solution is to identify the required constraints during plant design and to enforce them in the software and overall system design, including hardware and human behavior. Identifying these safety constraints is the goal of the new systemic hazard analysis techniques and the basis of the new systems approach to process safety.

Another example comes from the post mortem analysis of the events that led to the 2005 explosion of the isomerization unit at BP's Texas City refinery. The record indicated that there were malfunctioning sensors, stuck valves, and violation of operating procedures by operators, all of which can be seen as "component failures." The Baker Panel Report on this accident found that if one were to accept this tantalizingly attractive explanation of the accident, one would not have uncovered the systemic unsafe interactions at higher levels of management and overall operation of the plant in an unsafe manner, which led to these simultaneous component failures.

Safety and reliability are *different* system properties. *Reliability is a component property* and in engineering is defined as the probability that a component satisfies its specified behavioral requirements over time and under given conditions. Failure rates of individual components in chemical processes are fairly low, and the probability of simultaneous failure of two or more components is very low, unless these failures are not really independent as usually assumed.

In contrast, *process safety is a system property* and can be defined as absence of accidents, where an *accident* is defined as an event involving an unplanned and unacceptable loss. One does not imply nor require the other—a system can be reliable and unsafe or safe and unreliable. In some cases, these two system properties are conflicting, i.e., making the system safer may decrease reliability and enhancing reliability may decrease safety. For example, increasing the burst-pressure to working-pressure ratio of a tank often introduces new dangers of an explosion in the event of a rupture.

As chemical processes have become more economical to operate, their complexity has increased commensurably: There are many material recycles, heat and power integration, frequent changes of optimal operating points, integration of multivariable control and operational optimization. The type of accidents that result from *unsafe interactions* among the various process components is becoming the more frequent source of accidents.

In the past, the designs of chemical processes were intellectually manageable (serial processes with a few recycles, operating at fixed steady-states over long periods of time), and the potential interactions among the various system components could be thoroughly planned, understood, anticipated, and guarded against. In addition, thorough testing was possible and could be used to eliminate design errors before system use. Highly efficient modern chemical processes no longer satisfy these properties and system design errors are increasingly the cause of major accidents, even when all components have operated reliably, i.e. have not failed.

2. An Alternative System-Theoretic View of Process Safety

More effective process safety analysis methodologies and techniques, that avoid the limitations of those based on chain-of-events models, are possible if they are grounded on systems thinking and systems theory. Systems theory dates from the 1940s and 1950s and was a response to the limitations of the classic analysis techniques in coping with the increasingly complex systems being built after World War II.

In the traditional decomposition approach of classical chemical engineering, processing systems are broken into distinct unit operations and other operating components such as the elements of control loops. The behavior of the individual physical elements are modeled and analyzed separately and their behavior is decomposed into events over time. Then, the behavior of the whole system is described by the behavior of the individual elements and the interactions among these elements. A set of assumptions, however, underlies the reasonableness of this approach including:

(a) The separation of the process into its components is feasible, i.e. each component or subsystem operates independently and analysis results are not distorted when these components are considered separately.

(b) The processing components or behavioral events are not subject to feedback loops and non-linear interactions, and the behavior of the components is the same when examined singly as when they are playing their part in the whole.

(c) The principles governing the assembly of the components into the whole are straightforward, that is, the interactions among the subsystems are simple enough that they can be considered separate from the behavior of the subsystems themselves.

These assumptions no longer hold in modern petrochemical plant design. In contrast to the decomposition approach, the systems approach focuses on the processing system as a whole and does not decompose its behavior into events over time. It assumes that some properties of the processing system can only be treated adequately in their entirety, taking into account all facets related not only to the technical and physical-chemical underpinnings of the processing operations, but also the human, social, legislative and regulatory aspects surrounding the process itself. These system or *emergent* properties derive from the relationships among the parts of the system: how the parts interact and fit together. Thus, the system approach concentrates on the analysis and design of the whole as distinct from the components or its parts and provides a means for studying *emergent* system properties, such as process safety. A simple way to summarize the concept of emergent properties is the observation that "The whole is greater than the sum of its parts."

Using the system approach as a foundation, new types of accident analysis (both retroactive and proactive) can be devised that go beyond simply looking at events. They can identify the processes and systemic factors behind the losses and also the factors (reasons) for migration toward states of increasing risk. This information can be used to design controls that prevent hazardous states by changing the design of the processing system to prevent or control the hazards and state migration to higher-risk domains and, in operational systems, identify leading indicators that can detect the increasing risk before a loss occurs.

Preventing Failures vs. Enforcing Constraints on System Behavior

The traditional approach to process safety focuses on preventing component failure. In contrast, a systems approach emphasizes *enforcing constraints on system behavior*.

Process safety is not part of the mission or reason for the existence of a chemical plant. Its mission—its reason of existence—is to produce chemicals. To be safe in terms of not exposing bystanders and the surrounding environment to destructive effects of unleashed toxins or shock waves is a constraint on how the mission of a chemical plant can be achieved, where *by constraint we imply limitations on the behavioral degrees of freedom of the system components.* This seemingly trivial observation and statement has far reaching ramifications on how process safety should be viewed and how safe processes should be designed and operated. The reason for this assertion is simple:

> *Accidents occur when process-safety constraints are violated. Given that these constraints are imposed on the operational state of a chemical plant as a system, one concludes that process safety is a problem that must be addressed within the scope of an operating plant seen as a system.*

Accidents result from interactions among components that violate the safety constraints. These violations may result from inadequate monitoring of the safety constraints through absence or by providing inadequate control of behavior (both physical and human), which leads to insufficient corrective action.

> *A control-inspired view of process safety suggests that accidents occur when external disturbances, component failures, or unsafe interactions among processing components are not adequately handled by the existing control systems, leading to a violation of the underlying safety constraints.*

A chemical process does not operate as a purely engineered system, driven only by the physical and chemical or biological phenomena, and its safety cannot be viewed solely in terms of its technical components alone. Many other functions have a direct or indirect effect on how a process is operated, monitored for process safety, assessed for risk, and evaluated for compliance to a set of regulatory constraints (e.g. environmental constraints or process safety regulations). It is operated by human operators, it is serviced and repaired by maintenance personnel, and it is continuously upgraded and improved by process engineers. Operational managers of process unit areas or whole plants and managers of personal and process safety deploy and monitor process performance, execution of standard operating procedures, and compliance with health, safety, and environmental regulations.

Higher up in the hierarchy, company-wide groups define rules of expected behavior, compile best practices and safety policy standards, receive and analyze incident reports, and assess risks. Even higher in the hierarchy, local, state, or federal legislative and/or regulatory bodies define, deploy, and monitor the compliance of a set of rules, all of which are intended to protect the social and physical environment in which the process operates.

The idea of modeling socio-technical systems using process-control concepts is not a new one. Jay Forrester in the 1960s, for example, created *System Dynamics* using such an approach. Industrial engineering models often include both the management and technical aspects of systems. As one example, Johansson in 1985 described a production system as four subsystems: physical, human, information, and management. The physical subsystem includes the inanimate objects—equipment, facilities, and materials. The human subsystem controls the physical subsystem. The information subsystem provides flow and exchange of information that authorizes activity, guides effort, evaluates performance, and provides overall direction. The organizational and management subsystem establishes goals and objectives for the organization and its functional components, allocates authority and responsibility, and generally guides activities for the entire organization and its parts.

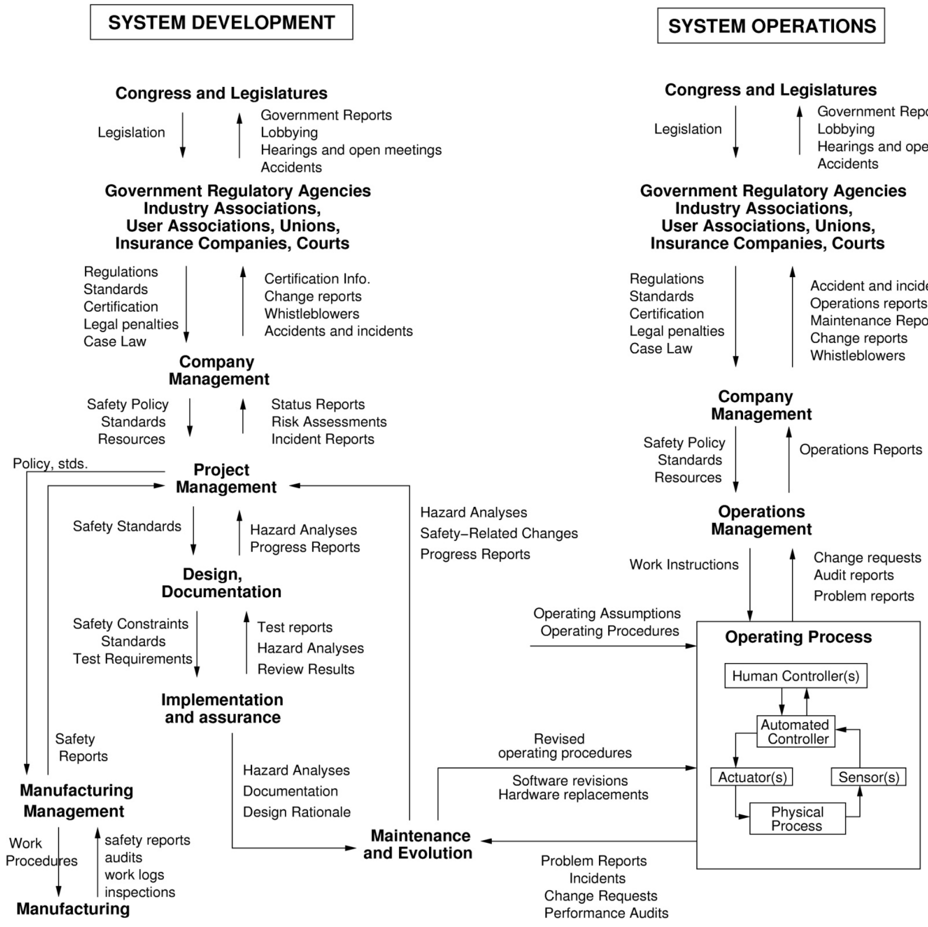

Figure 3. The hierarchical organization of control structures for the monitoring and control of safety constraints during the development and operation of a chemical process.

In a systems view, process safety is a control problem and is managed by a properly designed control structure, which is imbedded in an adaptive sociotechnical hierarchical system. An example is shown in Figure 3.

In Figure 3, the left-hand side of the diagram represents system development and the right-hand side shows operations. In both sides, the upward arrows represent the flow of monitoring and feedback while the downward arrows represent control.

Audits and reports on operations, accidents, and problem areas gathered at the lowest level of process operations provides important information for the adaptation of work instructions, maintenance and/or operating procedures, which in turn may lead to adaptation of company policies on safety and standards. In case of high profile destructive accidents, or a series of reported incidents, new tighter regulations may be imposed, or new legislation may be introduced to empower tighter regulations.

Methodologies on monitoring and diagnosing the operational state of a chemical process have received a lot of attention in chemical engineering. However, they have always been considered as self-standing systems, not part of a broader process safety program. Within the framework of the system-theoretic framework for process safety, monitoring and diagnosis tasks become essential elements of the adaptive control structure that ensures the satisfaction of the safety constraints. Integrating the two into one comprehensive system, that monitors the migration of operating states to higher risk domains, and diagnosing the reasons for this migration, becomes a natural task for process systems engineers. At the same time, leading indicators of increasing risk can be identified.

It is in the identification of the role of the entire control structure in accidents that the traditional hazard analysis techniques exhibit their most pronounced inadequacy: fault tree analysis, FMEA, bow-tie analysis, and HAZOP (or any of the other analysis techniques based on the event-chain model of accident causation) cannot identify all the pertinent safety constraints that arise from the unsafe interactions of processing components or from the interaction of management functions and processing operations. Instead, they focus on identifying potentially dangerous physical component failures. However, what we observe as behavior at the physical level of an operating plant has been decisively shaped by decisions made at higher levels.

3. Implementing a Systems View of Process Safety

The overall constraints on behavior must be enforced by both the design of the physical process and the design of the management and operations structures.

At the physical process level, using the prevailing chain-of-failure event causality models, the basic element is the failure of a component. In the control-inspired, system-theoretic view of process safety, the basic element is a safety constraint.

But, what exactly are the constraints? An obvious safety constraint in processing plants is the restriction that hydrocarbons to air ratio is outside the explosion range. Other technical constraints may involve restrictions on the pressures of vessels with gaseous components, levels of liquids in processing units, loads on compressors, operating temperatures in reactors or separators, or flows throughout the plant. Traditional operating envelops are manifestations of these safety constraints.

However, all of these constraints are local, that is, restricted to the operation of a single unit or a cluster of units with a common processing goal, such as the condenser and reboiler in conjunction to the associated distillation column. For example, material and energy balances in a dynamic setting should be obeyed at all time-points, and for all units, plant sections, and the entire plant. Monitoring these balances over time ensures that one will be able to observe the migration of operating states towards situations of high risk.

Emergent properties that relate the operations of processing units that are physically removed from each other are normally not constrained today because our simple event-chain accident causation models do not reveal such restrictions. The violation of material or energy balances over multi-unit sections of chemical plants is a typical example of an emergent property that is often overlooked.

Furthermore, emergent properties resulting from management functions are not constrained because they have never been the explicit goal of a comprehensive process safety treatment. For example, in the Texas City isomerization explosion, the repeated violation of safe operating procedures by the startup operators never constituted a violation of an important safety constraint in the minds of the supervisory management.

The interplay between human operators and safety constraints is crucial in ensuring the satisfaction of safety constraints. In times past, the operators were located close to process operations, and this proximity allowed a sensory perception of the status of the process safety constraints via direct "measurement", such as vibration, sound, and temperature. Displays were directly linked to the process via analog signals and thus were a direct extension of it. As the distance between the operators and the process grew larger, due to the introduction of electromechanical and then digital measurement, control, and display systems, the designers had to synthesize and provide a "virtual" image of the process operational state. The monitoring and control of safety constraints became more indirect, through the monitoring and control of the "virtual" process safety constraints.

Thus, modern computer-aided monitoring and control of process operations introduces a new set of safety constraints that the designers must account for. Designers must anticipate particular operating situations and provide the operators with sufficient information to allow the monitoring and control of critical safety constraints. This goal is possible using the system-theoretic view of a plant's operation, but it requires analysis and sophisticated design. The traditional analysis techniques, such as bow tie analysis, HAZOP, FTA, and FMEA, are not powerful enough to find them in today's more complex systems. For example, a designer should make certain that the information an operator receives about the status of a valve is related to the valve's status, i.e. open or closed, not on whether power had been applied to the valve or not, as happened in the Three Mile Island accident. Direct displays of the status of all safety constraints would be highly desirable, but present practices do not promise that it will be available any time soon.

In summary, accidents occur because safety constraints were never adequately imposed during process development and design, or, in the case that they were imposed, because they were inadequately monitored or controlled in the design or operation of the system. The solution is to identify the required constraints during plant design and to enforce them in the software and overall system design and operations. Identifying these safety constraints is the goal of the new systemic hazard analysis technique called STPA (System Theoretic Process Analysis).

How does STPA differ from the traditional analysis techniques used in the process industry? Instead of decomposing the process into its structural elements and defining accidents as the result of a chain of failure events, as prevailing techniques do, it describes processing systems and accidents in terms of a hierarchy of adaptive feedback control systems, as shown in Figure 3.

At each level of the hierarchy and for the processes within the physical operating plant (simplified in the lower right hand corner of Figure 3) there is a set of feedback control structures that ensure the satisfaction of the control objectives, i.e. of the safety constraints. Figure 4 shows the structure of a typical feedback control loop and the process models associated with the design of the controllers. In general, in each feedback loop the controller may be *automated* or may be a *human supervisor* (a human controller).

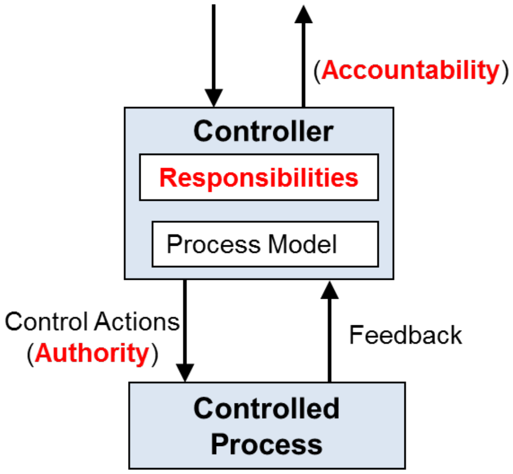

Figure 4. A typical but simplified control loop

Every controller must contain a model of the process it is controlling. In humans, we often call this model a "mental model." Whether the process models are embedded in the control logic of an automated controller or in the mental model maintained by a human controller, they must contain what the controller believes is the current state of the process and other information, such as the way that state can change or be changed. Accidents often occur when the controller's process model becomes inconsistent with the actual state of the process. At Texas City, for example, the operators thought that the level of liquid in the ISOM tower was below the safe level so they did not take steps to reduce the level of liquid. In fact, it was dangerously high. The operators (in hindsight) wrong behavior arose because of flaws in the system design (such as the wrong location of the high-level sensor], not in their procedures or behavior.

Furthermore, if, as is usually the case, human controllers (e.g. human operators) are controlling automated controllers, in addition to having a model of the controlled process, the human operator must have a model of the automated controller's behavior in order to monitor and supervise it.

In the chemical industry as well as most other industries, the number of accidents and near-misses caused by inaccurate process models is very high. They can be identified through the analysis procedures of STPA that generate the potential causal scenarios for accidents by examining how the process models could get into a state that differs from the actual state. These scenarios may be made up of the traditional component failures or parameter deviations that are considered in the standard hazard analysis methods used in chemical engineering, but they include many more scenarios than the causes that can be found by the standard methods, including component interactions, human mistakes, and engineering design errors. The scenarios derived from the analysis can then be used to improve the design or operation of the plant.

The model in Figure 4 is simplified. The new hazard analysis techniques, such as STPA (System Theoretic Process Analysis) and accident analysis techniques, such as CAST (Causal Analysis based on Systems Theory), use much more detailed and sophisticated models of the process including sensors, actuators, displays of the state of process operations (including posting of alarms), and the interfaces between all the automated and human control systems.

The STPA analysis is performed directly on the control model. HAZOP is also performed on a model, but that model is of the physical design of the plant. Such a physical model can be used to identify physical failures leading to accidents. STPA models are functional models and can be used to identify the functional design flaws as well as the physical failures that can lead to the violation of the safety constraints.

Despite our best efforts to prevent them, accidents will still occur. A systems approach to safety also requires more powerful accident investigation and analysis approaches that go beyond the superficial events preceding the loss and identify the systemic factors leading to the accident. CAST (Causal Analysis using Systems Theory) provides such an accident analysis tool.

Summary

By shifting the focus from component failures to violations of safety constraints, process safety engineering can be established within a broader and more comprehensive framework, fully supported by a new set of system-theoretic methodologies and techniques. As stated earlier, wide-spread use of STPA and the new systems approach in other industries is paying dividends, not only in identifying a larger set of causal factors but in support for much more effective accident prevention activity. Although STPA is more powerful than current hazard analysis techniques, it is actually less expensive because it approaches the problem in a different way.

The same system-theoretic approach also works for other important emergent system properties such as cybersecurity, productivity, and various aspects of quality although it is only beginning to be applied to these system properties.

Additional Reading: Nancy Leveson, Engineering a Safer World, MIT Press, 2012. There are also many papers and presentations at http://psas.scripts.mit.edu/home/