A New Approach to Risk Management and Safety Assurance of Digital Instrumentation and Control Systems

John Thomas (jthomas4@mit.edu)
Nancy Leveson (leveson@mit.edu)

Massachusetts Institute of Technology: 77 Mass Ave, Cambridge, MA 02139-4307

INTRODUCTION

One of the challenges of risk management in nuclear power plants is the rapid pace of change in technology, particularly the introduction of digital technology. Although digital instrumentation and control promises several benefits like self-checking, on-line diagnostics, improved accuracy, fault tolerance, and automated sensor calibration verification, it also presents unique challenges from software logic errors to unanticipated interactions that lead to unexpected or unsafe behavior [1].

The Problem

For traditional electro-mechanical safety systems, many effective methods exist for assurance and risk management. Unfortunately, many of the assumptions underlying these traditional methods do not apply to software. First, software does not fail randomly like hardware: software is pure design without any physical realization of that design—it therefore contains only design defects. Software can be thought of as design abstracted away from its physical representation, that is, it is pure design without any physical realization. While this abstraction reduces physical limits in design and thus allows exciting new features and functions to be incorporated in system design, it also greatly increases potential complexity and changes the types of failure modes.

Essentially there are two means for "failure" of digital systems. The hardware on which the software executes can fail in the same way that analog hardware does and the protection against these types of computer hardware failures, such as redundancy, is similar. In addition to the computer hardware failing, however, the software (which embodies the system functional design) can be incorrect or include behaviors that are unsafe in the encompassing system. Because the potential problems are always pure design defects, redundancy (which simply duplicates the design errors) is not effective. As Knight and Leveson have shown, making multiple versions of the software using different teams does not solve the problem either [2,3]. People make mistakes on the hard cases in the input space; they do not make mistakes in a random fashion. Therefore, even independently developed software modules are very likely to contain common cause failure modes.

In fact, almost all serious accidents caused by software have involved errors in the requirements, not in the implementation of those requirements in software code [4]. In most accidents, the software requirements have had missing cases or incorrect assumptions about the behavior of the system in which the software is operating. Often there is a misunderstanding by the engineers of the requirements for safe behavior, such as an omission of what to do in particular circumstances that are not anticipated or considered. The software may be "correct" in the sense that it successfully implements its requirements, but the requirements may be unsafe in terms of the specified behavior in the surrounding system, the requirements may be incomplete, or the software may exhibit unintended (and unsafe) behavior beyond what is specified in the requirements. Redundancy or even multiple versions of the implementations of the requirements does not help in these cases.

In addition, most software represents a new design—it is used to introduce efficiencies or functions that were not in previous hardware designs. Even reuse of old software does not seem to solve the problem [5,6]: almost all the software-related spacecraft losses in the past few decades involved reused software from past spacecraft [7]. These results may stem partly from complacency created by successful use in previous systems and because undocumented assumptions made during the original development may be inappropriate for the new use.

The violation of these basic causal assumptions about accidents means that many of the traditional techniques for safety assurance do not apply to the digital components of systems. The problem then is how can we improve our ability to provide software assurance for safety-critical applications and how can these techniques be combined with traditional design and assurance techniques to provide a more effective means of designing and evaluating mixed analog and digital instrumentation in NPPs.

A Potential Solution

To address these problems, a new accident causality model called STAMP (System-Theoretic Accident Model and Processes) was created based on system theory to include a broader view of accident causation and indirect or non-linear interactions among events [6]. In STAMP, safety is reformulated as a control problem rather than simply a reliability (or availability) problem. Component failure (and unreliability of the system components) is still included, but more generally as accidents may occur when component failures, external disturbances, or unsafe interactions among system components are not adequately

handled, i.e., controlled, resulting in unsafe system behavior.

In STAMP, in order to provide effective control the controller must have an accurate model of the process it is controlling (Figure 1). For human controllers, this model is commonly called the mental model. Regardless of whether the controller is automated or human, the process or mental model is used by the controller to determine what control actions are necessary to keep the system operating effectively. The process model includes assumptions about how the controlled process operates and about the current state of the controlled process. Accidents in complex systems, particularly those related to software, often result from inconsistencies between the model of the process used by the controller and the actual process state, which leads to the controller providing unsafe control.

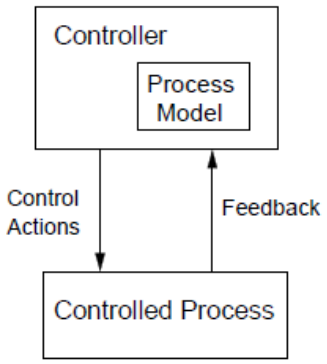

Figure 1: Controller Process Models.

Using these concepts, we have created a new hazard analysis technique called STPA (System Theoretic Process Analysis) [6]. STPA can be used to identify the safety constraints that must be enforced and to ensure that the system design adequately enforces them. It also identifies the required process model (mental model if the controller is a human) that the controller needs in order to provide adequate control and thus the information required in that process or mental model. If that information gets lost or corrupted, accidents can occur.

STPA is a method for examining the control loops in the safety control structure to find potential flaws and the potential for (and causes of) inadequate control. In this framework, there are four types of inadequate control that can lead to accidents:

1. Incorrect or unsafe control commands are given
2. Required control actions (for safety) are not provided
3. Potentially correct control commands are provided at the wrong time (too early or too late), or
4. Control is stopped too soon or applied too long

Once the potential for inadequate control is identified, constraints can be formed and the causes of inadequate control (e.g. process model flaws, incorrect or missing feedback, etc.) can be identified. Figure 2 shows the generic control flaws and inadequate control that can be identified using STPA.

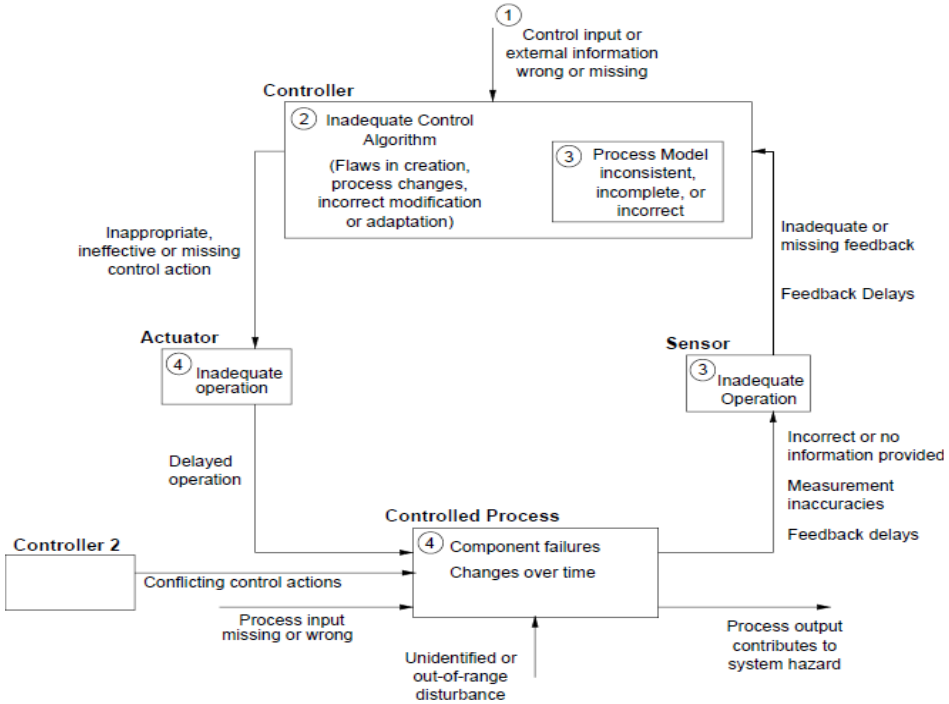

Figure 2: A classification of control flaws leading to hazards

Because this framework extends current accident models and thus includes component failure accidents, STPA not only identifies the hazard scenarios identified by fault trees, event trees, and other traditional hazard analysis methods, but it also includes those factors not included or poorly handled in these traditional methods such as software requirements errors, component interaction accidents, complex human decision-making errors, inadequate coordination among multiple controllers, and management and regulatory decision making.

DESCRIPTION OF THE WORK

Our research goal was to evaluate the applicability, feasibility, and relative efficacy of using STPA in digital nuclear power plants. The case study for this research involves a generalized version of an EPR (Evolutionary Power Reactor), a version of which appears in [8]. The EPR studied is a type of PWR (Pressurized Water Reactor). The system includes one Steam Generator (SG) and one Main Steam Isolation Valve (MSIV). The EPR reactor is fully digital, that is, all control systems, including the Reactor Protection System, are digital. The analysis focuses on a sub-set of the Nuclear Power Plant (NPP) system: the systems involved in closing the Main Steam Isolation Valve (MSIV). The same process could be applied to the rest of the system.

System Overview

A generic diagram of a PWR is shown in Figure 3. During normal operation, the coolant in the primary cooling system (left of the diagram) transfers heat from the reactor to the Steam Generator (SG). The SG contains water that cools the primary coolant and evaporates into steam. The SG prevents primary coolant, which is radioactive, from mixing with the water, which is not radioactive. The steam produced in the SG travels to a turbine connected to a generator to produce electricity. The steam is cooled in the condenser and pumped back into the SG to begin the cycle again. The loop formed by the SG, turbine, and condenser is known as the secondary cooling system.

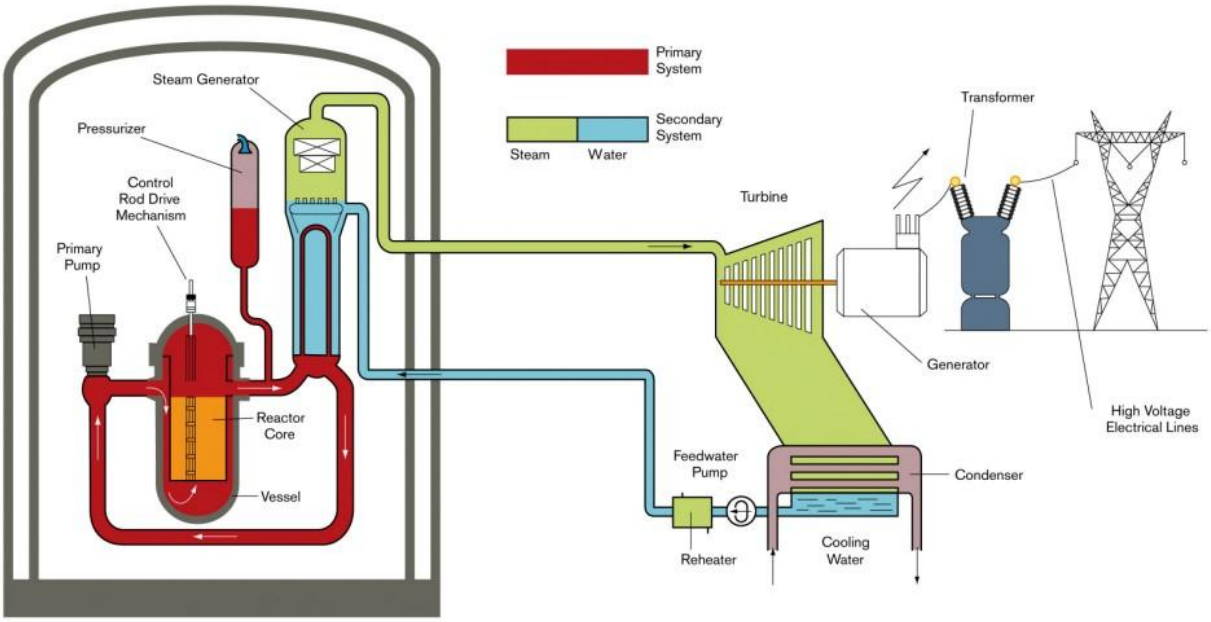

Figure 3: Pressurized Water Reactor (Diagram from [9])

The MSIV is a valve located on the main steam line from the SG. During normal operation, the MSIV is kept open to permit cooling of the primary cooling system via the secondary system. In case of an abnormal situation, the MSIV can be closed to isolate the SG from the rest of the secondary system. MSIV closure is necessary if there is a break in the main feedwater pipe to the SG that allows water to leak out, an internal SG Tube Rupture (SGTR) that allows primary coolant to mix with secondary water, or a break in the main steam line exiting the SG.

Because MSIV closure prevents the secondary system from adequately cooling the primary system, a number of backup systems are provided to cool the primary coolant in case of MSIV closure. These backup systems may include redundant SGs, turbine bypass valves, main steam relief isolation valves (MSRIV) and main steam relief control valves (MSRCV), safety relief

valves (SRV), the Chemical Volume Control System (CVCS), and the Emergency Core Cooling System (ECCS). These systems are included in the analysis only to the extent that they impact the decision to close the MSIV.

The STPA analysis that follows begins by identifying the accidents, hazards, and control structure for the overall system. The remaining steps analyze in more detail the systems related to closure of the MSIV.

System-level Accidents

The first step in STPA is to identify the system-level losses, or accidents, to be considered. Accidents often involve loss of human life or injury, but any loss can be included that is unacceptable and must be prevented. Table 1 below shows the system-level accidents that are analyzed in this analysis.

Table 1: NPP system-level accidents to be prevented

A-1: People injured or killed
A-2: Environment contaminated
A-3: Equipment damage (economic loss)
A-4: Loss of electrical power generation

People injured or killed (A-1) includes both employees and the general population, and may involve radiation exposure, explosion, or any other mechanism. Environment contaminated (A-2) includes radiation or other harmful release to the air, ground, or groundwater, or any other part of the environment. Equipment damage (A-3) refers to the economic loss associated with any damage to equipment regardless of whether any radiation is released. Loss of electrical power generation (A-4) includes any unplanned plant shutdown.

Priorities may be assigned as not all accidents are equally important. In addition, the accidents need not be mutually exclusive, and in fact it is possible to experience all four losses at once. Finally, economic damage such as equipment loss or the loss of electrical power generation (A-4) may not be of immediate importance in a licensing review or a traditional safety analysis but it is certainly a concern for the utility. STPA can be used for any type of loss that is important to those doing the analysis. Incorporating other types of losses, such as mission or economic losses, can not only allow better decision making with respect to achieving multiple requirements but can also assist in identifying and making tradeoffs between conflicting goals.

System-level Hazards

Once the system accidents have been defined, the hazards can be identified. Table 2 summarizes the hazards included in this analysis and the accidents to which they are related.

Table 2: NPP system-level hazards

Hazard	**Related Accident**
H-1: Release of radioactive materials	A-1, A-2
H-2: Reactor temperature too high	A-1, A-2, A-3, A-4
H-3: Equipment operated beyond limits	A-3, A-4
H-4: Reactor shut down	A-4

Release of radioactive materials (H-1) refers to any release outside the primary system, including releases into the secondary cooling system, groundwater, and air inside or outside the containment structure(s). These releases must be controlled to prevent exposure to people or the environment (A-1 and A-2). Reactor temperature too high (H-2) is a dangerous condition that can cause every system-level accident (for example, if the fuel rods melt), or it may lead to A-1 and A-2 without any radiation release (for example, through hydrogen production or other dangerous conditions). Although H-2 may exist without an accident (for example, if there is a hydrogen explosion but containment holds), H-2 is a dangerous condition that should be controlled in the design. Equipment operated beyond limits (H-3) includes operation beyond safe limits that causes reactor damage or operation beyond design limits that causes damage to other equipment. Reactor shut down (H-4) includes any unplanned shutdown that may result in a loss of electrical power generation.

Safety Control Structure

The next step in STPA is to develop the safety control structure, including the controllers and the functional control/feedback paths in the architecture. The high-level control structure for this system is shown in Figure 4. The components inside the dashed box control the closing of the MSIV. They are analyzed in further detail for the remainder of the case study. Figure 5 shows a more detailed control structure for the systems highlighted in the dashed box.

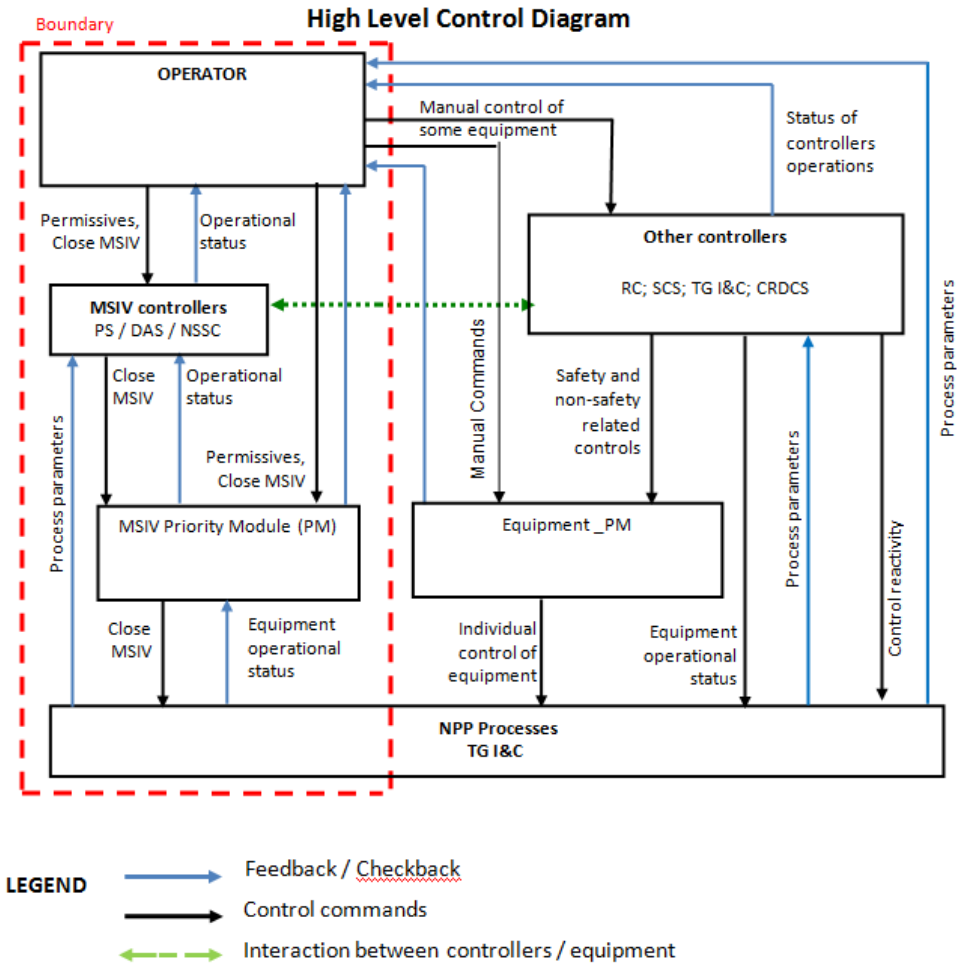

Figure 4: High-level PWR Safety Control Structure

The dotted arrow represents the communication between the MSIV controllers and other controllers. For example, the Protection System (PS) contacts the Safety Control System (SCS) in order to initiate the Engineering Safety Features (ESF) controls following ESF actuation. The Reactor Controls (RC) controller also communicates with Non-Safety System Controller (NSSC) in order to provide command signals for actuators used in RC functions other than control rods, such as the BMC (Boron and Makeup Control) components for Boron control.

There are four controllers that can provide a control action to close the MSIV: the Operator, the NSSC, the PS, and the Diverse Automation System (DAS). These four controllers send control actions to the MSIV Priority Module (PM), which uses a pre-programmed priority setting to determine which control actions to forward to the MSIV actuator. In this sense, the PM can also send control actions.

If the operator detects a need to close the MSIV, he or she may issue a Close MSIV command to the PM. The PM determines which controller is in charge according to a priority scheme, and forwards commands directly to the MSIV actuator. In this case, the PM would normally forward the command from the operator to the MSIV actuator.

The operator may also send a Close MSIV command to the NSSC, which provides manual control for the MSIV. In this situation, the NSSC would normally forward the command from the operator to the PM, which would then forward the command to the MSIV actuator.

The PS is an automated system that can automatically detect some situations in which a Close MSIV command is necessary. In these situations the PS can provide the Close MSIV command to the PM which can forward the command to the MSIV actuator.

Finally, the DAS is a backup protection system that is used if there is a problem with the PS. The DAS can issue a Close MSIV command to the PM, which would normally forward the command to the MSIV actuator.

Figure 5 shows the control structure of interest in more detail, and the main process or mental model variables related to MSIV closure are listed inside the controllers.

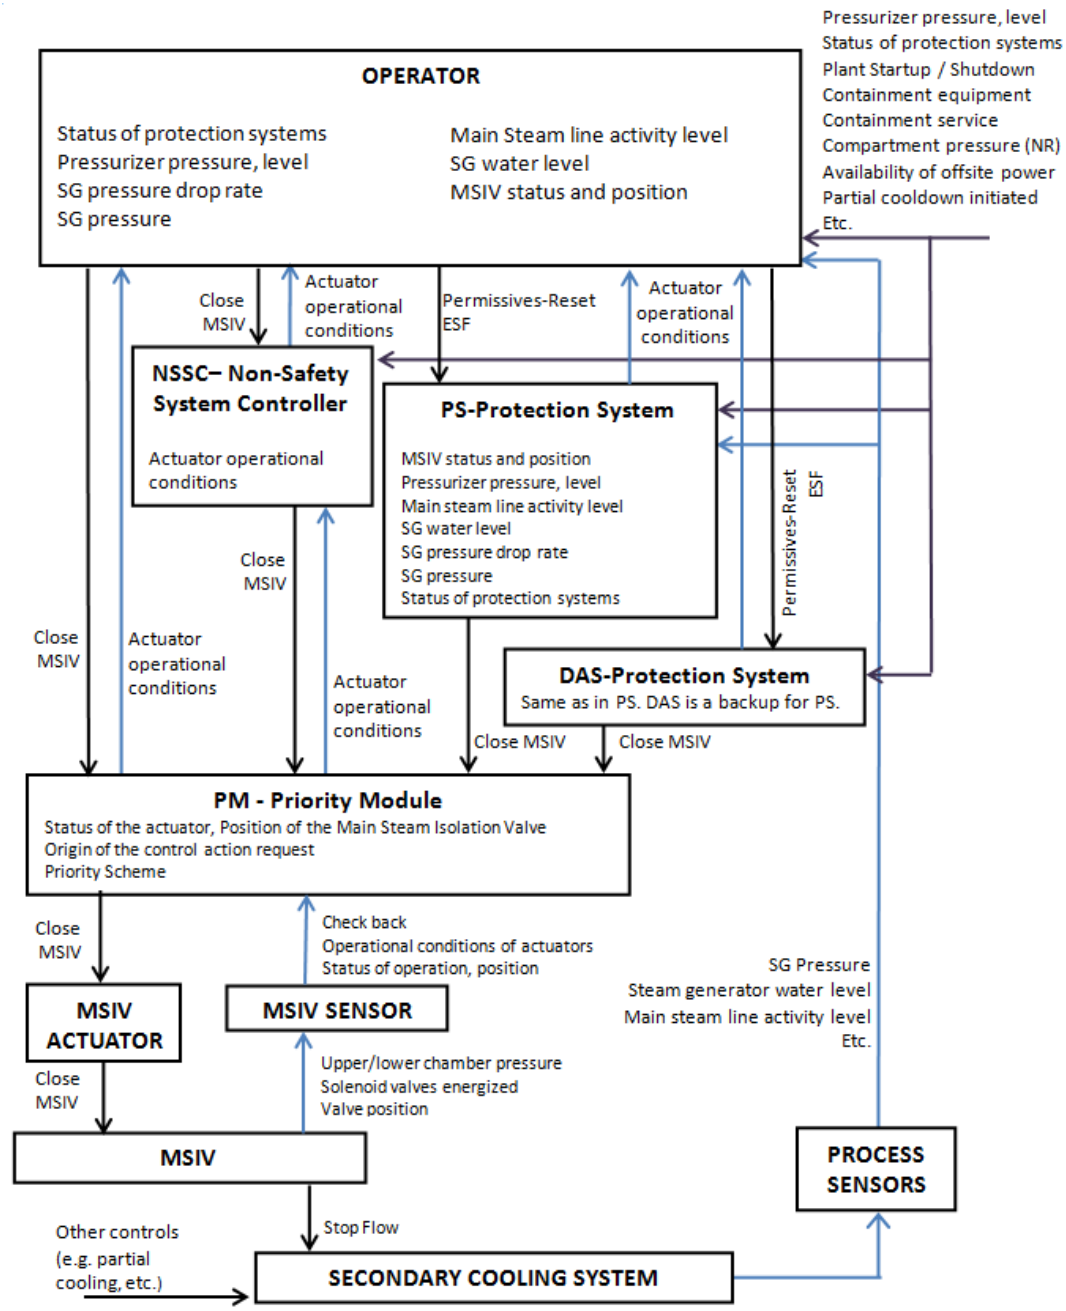

Figure 5: More detailed safety control structure for MSIV closure

Hazardous Control Actions

Once the control structure has been developed, the analysis proceeds to identify unsafe control actions that can lead to the identified hazards. When considering potentially hazardous control actions, it is important to avoid assuming that other defense barriers are intact or that they are appropriate, sufficient, and error-free. For example, even if there is an emergency feedwater system to provide the necessary cooling in the event of a relief valve inadvertently commanded open, it is still hazardous to inadvertently command the relief valve open. These hazardous actions must be included in the analysis and prevented regardless of other protective systems intended to mitigate unsafe behavior.

Table 3 summarizes the hazardous control actions that were identified for the control action command *Close MSIV*.

Table 3: Hazardous Control Actions for Close MSIV

Control Action	Hazardous Control Actions			
	Not Providing Causes Hazard	**Providing Causes Hazard**	**Wrong Timing or Order Causes Hazard**	**Stopped Too Soon or Applied Too Long**
Close MSIV	Close MSIV not provided when there is a rupture in the SG tube, leak in main feedwater, or leak in main steam line [H-2, H-1, H-3]	Close MSIV provided when there is no rupture or leak [H-4] Close MSIV provided when there is a rupture or leak while other support systems are inadequate [H-1, H-2, H-3]	Close MSIV provided too early (while SG pressure is high): SG pressure may rise, trigger relief valve, abrupt steam expansion [H-2, H-3] Close MSIV provided too late after SGTR: contaminated coolant released into secondary loop, loss of primary coolant through secondary system [H-1, H-2, H-3] Close MSIV provided too late after main feedwater or main steam line leak [H-1, H-2, H-3, H-4]	N/A

Although hazardous control actions can be identified in an ad-hoc manner, a systematic approach has been developed as part of this research to provide additional guidance and to help identify unanticipated or off-nominal contexts. This systematic approach was used to identify the hazardous control actions in Table 3.

The systematic approach is based on the recognition that the safety of a control action is intimately connected to the context in which it is provided. For example, is it hazardous to open the MSIV? The answer depends on the context—the system state or state of the environment in which the command is given. The systematic procedure involves identifying potential control actions, identifying potentially hazardous contexts, and then analyzing which combinations together yield a hazardous control action.

The procedure can be divided into two parts, and each part can be performed independently of the other. The first part deals with control actions that are provided under conditions that make the action hazardous. The second part deals with control actions that are *not* provided under conditions that make *inaction* hazardous.

To identify control actions that are hazardous to provide, the context table in Table 4 was constructed. The columns on the left represent high-level process model variables and can be identified from the system hazards and the control structure. Each row represents a unique context that the control action may be provided in. The columns on the right specify whether the row represents a valid hazardous control action.

Column 6 in Table 4 specifies in which contexts it is hazardous to provide the *Close MSIV* control action. For example, row 1 describes a situation in which it is hazardous to close the MSIV: if there is no SG tube rupture, no main feedwater pipe leak, and no main steam line leak, then there is no need to close the MSIV. Closing the MSIV will cause H-4 (reactor shut down). If the operation of other support systems cannot make up for the additional heat exchange required, closing the MSIV will also lead to a loss of necessary cooling (H-2 in row 9 column 6).

If other support systems, including other CVCS, SI, ECCS, etc., are producing the additional cooling required during a rupture/leak, then closing the MSIV is not hazardous (rows 2-8, column 6) and a reactor shutdown is initiated regardless of any MSIV actions. If for some reason the other systems are not capable of producing the additional cooling needed, then closing the MSIV may cause other hazards (rows 10-16, column 6) including excessive temperature increase (H-2), release of radioactive materials (H-1), an immediate reactor shutdown or SCRAM (H-4) if not already triggered, and additional equipment damage (H-3). Depending on the type of rupture, it may actually be better to keep the MSIV open to control the temperature of the reactor (H-2) even though that would permit some radioactive steam to be introduced into the secondary system (H-1).

Table 4: Context table for *Close MSIV* control action provided

	1	2	3	4	5	6	7	8
	Control Action	**Steam Generator Tube**	**Condition of Main Feedwater Pipe**	**Condition of Main Steamline**	**Operation of other support systems**	**Control Action Hazardous?**	**Control Action Hazardous if Too Late?**	**Control Action Hazardous if Too Early?**
1	*Close MSIV*	Not Ruptured	No Leak	No Leak	Adequate	H-4	H-4	H-4
2		Ruptured	No Leak	No Leak	Adequate	No	H-1, H-2, H-3, H-4	H-3, H-4
3		Not Ruptured	Leak	No Leak	Adequate	No	H-2, H-3, H-4	No
4		Not Ruptured	No Leak	Leak	Adequate	No	H-2, H-3, H-4	No
5		Ruptured	Leak	No Leak	Adequate	No	H-1, H-2, H-3, H-4	H-3, H-4
6		Not Ruptured	Leak	Leak	Adequate	No	H-2, H-3, H-4	No
7		Ruptured	No Leak	Leak	Adequate	No	H-1, H-2, H-3, H-4	H-3, H-4
8		Ruptured	Leak	Leak	Adequate	No	H-1, H-2, H-3, H-4	H-3, H-4
9		Not Ruptured	No Leak	No Leak	Inadequate	H-2, H-4	H-2, H-4	H-2, H-4
10		Ruptured	No Leak	No Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4
11		Not Ruptured	Leak	No Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4
12		Not Ruptured	No Leak	Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4
13		Ruptured	Leak	No Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4
14		Not Ruptured	Leak	Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4
15		Ruptured	No Leak	Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4
16		Ruptured	Leak	Leak	Inadequate	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4	H-1, H-2, H-3, H-4

Although providing a control action can be hazardous, *not* providing a control action can be equally hazardous. Table 5 shows the context table for *not* providing the *Close MSIV* control action. As before, a reactor shutdown should be initiated for any rupture regardless of the MSIV control action. However because these tables are used to identify hazardous control actions, only hazards that are affected by an absent *Close MSIV* control action are listed at this stage of the analysis.

Table 5: Context table for *Operator does not provide Close MSIV control action*

	1	2	3	4	5	6
	Control Action	**Steam Generator Tube**	**Condition of Main Feedwater Pipe**	**Condition of Main Steamline**	**Operation of other support systems**	**Not providing control action is hazardous?**
1	*Close MSIV*	Not Ruptured	No Leak	No Leak	Adequate	No
2		Ruptured	No Leak	No Leak	Adequate	H-1, H-2, H-3, H-4
3		Not Ruptured	Leak	No Leak	Adequate	H-2, H-3
4		Not Ruptured	No Leak	Leak	Adequate	H-2, H-3
5		Ruptured	Leak	No Leak	Adequate	H-1, H-2, H-3, H-4
6		Not Ruptured	Leak	Leak	Adequate	H-2, H-3
7		Ruptured	No Leak	Leak	Adequate	H-1, H-2, H-3, H-4
8		Ruptured	Leak	Leak	Adequate	H-1, H-2, H-3, H-4
9		Not Ruptured	No Leak	No Leak	Adequate	No
10		Ruptured	No Leak	No Leak	Inadequate	H-1, H-2, H-3, H-4
11		Not Ruptured	Leak	No Leak	Inadequate	H-2, H-3
12		Not Ruptured	No Leak	Leak	Inadequate	H-2, H-3
13		Ruptured	Leak	No Leak	Inadequate	H-1, H-2, H-3, H-4
14		Not Ruptured	Leak	Leak	Inadequate	H-2, H-3
15		Ruptured	No Leak	Leak	Inadequate	H-1, H-2, H-3, H-4
16		Ruptured	Leak	Leak	Inadequate	H-1, H-2, H-3, H-4

A comparison of Table 4 and Table 5 shows that there are conflicts that must be resolved. In both tables, rows 10 to 16 are marked as hazardous. In other words, in these situations it is hazardous to close the MSIV yet hazardous to keep the MSIV open. In some cases, it is possible to revisit the design to eliminate the conflict and provide a safe option. If the conflict cannot be resolved, a decision must be made about what action should be taken in these contexts, that is, which is the *least* hazardous? For this case study, after consultation with nuclear engineers and regulators it was found that rows 10 to 16 may not have been analyzed in previous safety analyses with respect to MSIV control. For the purposes of this research, the consensus was to assume that it may be best to keep the MSIV open in the context of row 10 to maximize the amount of cooling provided even though doing so will contaminate the secondary cooling system and eventually require costly repairs. Rows 11-16, on the other hand, involve leaks in the pipe supplying water to the steam generator and/or the line that carries steam away. If the MSIV is left open in these situations, the amount of water in the steam generator can decrease and eventually lead to less cooling capability or an overcooling transient. Therefore, in these situations (rows 11-16), it was assumed that it may be best to keep the MSIV closed to maximize the amount of cooling provided even though it is only a temporary measure. These solutions were found to differ from current designs of MSIV controllers, which do not act based on the state of other support systems and may automatically close the MSIV during any rupture.

Both of these assumptions should be reviewed and evaluated carefully by domain experts. The purpose of this research case study was not to provide final solutions to these hazardous situations, but to develop and apply hazard analysis methods that can uncover hazardous control and provide the safety-critical questions that need to be considered. Note that although Table 4 and 5 use high-level contexts, the analysis can also be performed in more detail using the lower level process model variables in Figure 5. A more detailed analysis could be necessary if, for example, it is found that the best solution depends on the type of steam generator tube rupture, the amount of pressure in the SG, etc.

Of course, in any of these situations, there are other control actions that need to take place outside the MSIV control loop—they can be analyzed using the same approach. In addition, every effort should be made to prevent many of these contextual conditions from existing in the first place. However, that is no excuse for incomplete requirements—correct behavior needs to be specified even for off-nominal situations.

The resulting hazardous control actions can then be summarized and converted into safety constraints as shown in Table 6.

Table 6: Unsafe Control Actions and Safety Constraints

Unsafe Control Action	**Safety Constraint**
UCA 1: Close MSIV not provided when there is a leak (rupture in the SG tube, leak in main feedwater, or leak in main steam line) and the support systems are adequate	**SC 1**: MSIV must be closed when there is a leak (rupture in the SG tube, leak in main feedwater, or leak in main steam line) and the support systems are adequate
UCA 2: Close MSIV not provided when there is a main feedwater or main steam line leak and other support systems are inadequate	**SC 2**: MSIV must be closed when there is a main feedwater or main steam line leak and other support systems are inadequate
UCA 3: Close MSIV provided when there is a SGTR but support systems are inadequate	**SC 3**: MSIV must not be closed when there is a SGTR and support systems are inadequate
UCA 4: Close MSIV provided too early (while SG pressure is high)	**SC 4**: MSIV must not be closed too early while SG pressure is too high
UCA 5: Close MSIV provided too late after rupture/leak (in the SG tube, main feedwater, or main steam line)	**SC 5**: MSIV must not be closed too late after rupture/leak (in the SG tube, main feedwater, or main steam line)
UCA 6: Close MSIV provided when there is no rupture/leak	**SC 6**: MSIV must not be closed when there is no rupture/leak

Causal Factors

The last part of STPA identifies potential causes of unsafe control actions and safety constraint violations including controller process model flaws, flawed requirements, design flaws, component failures, insufficient feedback, and other causes. This step is guided by the STPA classification of control flaws shown in Figure 2. The full results from this part of the analysis cannot be shown here due to space limitations, but a portion of the results appears below:

UCA 1: Close MSIV command not provided when there is a leak (rupture in the SG tube, leak in main feedwater, or leak in main steam line) and the support systems are adequate.

(1) Secondary cooling system (CVCS or emergency feedwater system)

a. Concurrent situation masks another. For example, a feedwater problem could happen concurrent with a SGTR causing the SG water level to stay practically stable, or NSSC engages PZR heaters to make up for loss of RCS pressure during SGTR.
b. Sensor component failures
c. ...

(2) Operator

a. Operator believes Steam Generator is not ruptured when it is ruptured
b. Operator believes the main steam line has no leak when it has a leak
c. Operator believes the main feedwater has no leak when it has a leak

d. Operator waits for the PS to handle the situation (e.g. Operator recognizes possible SGTR but believes PS will handle it)
e. Operator uncertainty regarding rupture/leak (conflict between being conservative under uncertainty versus immediate manual spurious shutdown which costs money and may be discouraged. May also prefer to wait for the automated system to resolve the problem versus intervening under uncertainty)
f. …

CONCLUSIONS AND FURTHER WORK

This work extended STPA to include more systematic methods for identifying and evaluating hazardous control actions and safety constraints, and demonstrated that it is both practical and feasible on digital NPP systems. Several potentially unsafe behaviors were identified and flagged for further consideration. Other research based on this work has developed advanced methods and automated processes that can further reduce the effort required to apply STPA to digital NPP systems [10].

Although the full analysis could not be included here due to space limitations, the basic process was demonstrated and several important insights can be derived. An example insight obtained from the analysis is the difficulty of detecting a Steam Generator Tube Rupture (SGTR) through the normal indicators, which can lead to a delayed response by the automated controllers and the operator. Current solutions rely on (i.e., give credit to) the operator's ability to detect and intervene in certain cases. Relying on the operator, however, may not be necessary or effective because of other factors that will influence the operator decision-making process. These factors are identified in STPA as possible causes for the operator not to provide the control action to close the MSIV or to provide it too late. The identified factors can be used to improve the design to make the operator error less likely or to mitigate it.

One reasonable recommendation, for example, is for the designers to simplify the indicators for the case of SGTR by making the level of radiation at the Main Steam Line a major indication to isolate the affected SG. This way, the automated Protection System (PS) would be able to detect the event earlier. In the current design, an indication of radioactivity is not sufficient for the PS to take action, and, as a result, there are additional scenarios in which neither the operator nor the PS may take action. For example, the operator may feel pressed to avoid spurious shutdowns and, as a consequence, he or she may wait longer for stronger evidence of the real problem. This type of response, in fact, is a common one identified by human factors experts in many real accidents. There could also be a situation where, after many years of work, the operator learns to completely rely on the automated controls to handle some incidents and becomes overconfident in its correct operation. This overreliance could lead to non-action or delayed action even though existing analyses have assumed he or she will immediately take action in that case.

The introduction of digital systems exacerbates the problem. Software allows highly complex systems to be created. While identifying safety-critical versus non-safety-critical components in a nuclear power plant was relatively straightforward for primarily electromechanical designs, the extensive use of software allows much more complex designs than previously possible and the potential for unintended and unexpected interactions among components. The more interactions between system components and the more complex the functional design, the more the opportunities for unintended effects and, consequently, the more opportunities for unsafe control actions that can lead to hazards. For example, in this digital NPP the operator has to manually change settings by manipulating priority logic in order to allow the NSSC to process the manual commands. This requirement can be a problem in case of an emergency.

The STPA analysis in this case study was limited in scope to the MSIV control and publically available information, but a more detailed STPA analysis seems warranted due to the central importance of this equipment in the control system. These are only some of the flaws or weaknesses in the design that can be identified from the partial system modeling and STPA hazard analysis performed for this limited research effort. Further STPA analysis results can be found in [8], and a more complete modeling and analysis effort would most likely uncover even more.

REFERENCES

[1] NEI, "USA's first fully digital station," 2011.
[2] Knight, J. C. and Leveson, N. G. "An experimental evaluation of the assumption of independence in multiversion programming," IEEE Trans. on Software Engineering, Jan. 1986
[3] Knight, J. C. and Leveson, N. G. "A reply to the criticisms of the Knight & Leveson experiment," SIGSOFT Software Engineering Notes, Jan. 1990
[4] Leveson N., *Safeware*, Addison-Wesley, 1995.
[5] Joyce J., "Software Safety for Air Traffic Management Systems," 21rst Digital Avionics Systems Conference, October 2002, IEEE Proceedings
[6] Leveson N., *Engineering a Safer World*, MIT Press, 2012.
[7] Leveson, N.G. "The Role of Software in Spacecraft Accidents," AIAA Journal of Spacecraft and Rockets, 2004.
[8] Thomas, J., F. Lemos, and N. Leveson, "Evaluating the Safety of Digital Instrumentation and Control Systems in Nuclear Power Plants", NRC Technical Research Report 2013.
[9] Syson, D., "The £3bn child of Chernobyl", in Daily Mail 2008
[10] Thomas, J. "Extending and Automating a Systems-Theoretic Hazard Analysis for Requirements Generation and Analysis," MIT PhD Dissertation, 2013