Visible Language 43.2/3, November 2009
Special Issue on Design Communication Error
pp. 170-195

LESSONS FROM THREE MILE ISLAND: VISUAL INFORMATION DESIGN IN A HIGH-STAKES ENVIRONMENT

AXEL ROESLER

University of Washington
Roesler, 170-197

© Visible Language, 2009
Rhode Island School of Design
Providence, Rhode Island 02903

ABSTRACT

Complex systems with mediated control at a distance are explored using the Three Mile Island nuclear accident of 1979 as the focus. In such a high-stakes environment, representations of operations are critical to support human-machine interactions and monitor safe operations. A time-line of the critical first minutes of the event is presented and an analysis of operations in the control room from a communication perspective point toward principles for a better design. While the case of Three Mile Island is well documented from an engineering perspective, its relationship to communication design and interaction design provide insight with regard to necessary collaboration across disciplines.

The major difference between a thing that might go wrong and a thing that cannot possibly go wrong is that when a thing that cannot possibly go wrong goes wrong it usually turns out to be impossible to get at or repair.

Douglas Adams, Mostly Harmless – Hitchhiker's Guide to the Galaxy Book 5 (1992)

At 4:00:37 a.m. on March 28, 1979, the Three Mile Island nuclear power plant near Harrisburg, Pennsylvania, encountered problems in the reactor cooling system that initiated the first, and to this day most severe, nuclear accident on U.S. soil. In its course, the second reactor of the plant (TMI-2) underwent a loss of coolant accident that resulted in severe reactor damage and the release of radioactivity into the environment. The release of radioactive material and threat of a nuclear catastrophe drew large-scale public alarm and political attention. Irreversible damage and severe contamination of the plant led to significant clean-up costs. Public protest and the ensuing discussions about safety, open communication and emergency response strategies made the Three Mile Island accident a significant event in the debate on the future of nuclear power in the U.S. The nuclear power industry, regulatory boards, emergency response authorities, the state of Pennsylvania and the nation all found themselves confronted with a technological accident of scale and public impact previously unknown. Extensive accident reports (Kemeny, 1979; Rogovin, 1980) prepared by the government and the Nuclear Regulatory Commission, along with various long-term studies of the effects of the accident, provide multiple, detailed perspectives on its causes, consequences and impact (Walker, 2004).

During investigations in the aftermath of the accident, the design of the control room was identified as a significant cause of the accident. The design and arrangement of information displays and controls in the control room confused operators about what was going on and negatively affected decision-making during critical phases of the accident.

Better integrated visual displays of plant status, safety margins and trends of operations could have prevented this severe nuclear accident which permanently damaged a $900 million nuclear reactor, caused nuclear contamination clean-up costs of nearly one billion dollars and potentially could have harmed thousands of people.

Thirty years later, findings from the examination of the control room are still as important as they were when an international community of human factors specialists and systems engineers were surprised by the magnitude of the accident (Woods et. al., 1994). Insights gained from the accident at Three Mile Island have led to better design in the control room systems that followed it, preventing similar breakdowns (with perhaps even larger consequences).

The exploration of the Three Mile Island accident reveals how interaction design and information display in the control room broke down during the dynamic development of the accident when many events took place both sequentially and in parallel at a very fast pace. Poor information representation played a key role in decision making of the control room operators during their response to the many problems that confronted them. The problems emerged rapidly during the initial five minutes of the accident beginning at 4:00 a.m. and remained for the next 150 minutes during the development of the accident. The patterns of interface failure that occurred during these initial hours of the accident are representative for design challenges in high-stakes work settings where technology supports the work of expert practitioners. Besides nuclear power plants, the observed patterns apply to the design of human-computer interactions in related high-stakes environments such as chemical processing plants, flight deck operations in aviation, the control of space systems and the management of anesthesia during surgery.

The successful design of interactive systems in these high-stakes domains relies on a synthesis of design principles in a concurrent, interdisciplinary design process that converges a) the design of human-computer interactions with insight from research on the complex behavior of technological systems and b) the evaluation of process control systems in work settings and the study of patterns that govern the cognitive aspects of work supported by (computerized) machines. The conditions of work are driven by human factors, the design of artifacts, procedures and organizations that determine training, operations, regulations, maintenance and licensing in the work domain.

When the TMI-2 control room failed to provide effective information representations that guided the response interactions of control room operators three decades ago, researchers and designers realized that they needed better models for the relationship between humans and machines in control rooms. This initiated new types of studies on the relationship between visual information display, human reasoning, interventions and the constraints that technical systems impose on sense-making strategies—as a result previously engineering-centered approaches to control room design began to open more towards a

Figure 1
The Three Mile Island-2 control room on April 3, 1979 (National Archives photo no. 220-TMI-DE9040061-13).

multidisciplinary approach. Design teams today involve human factors researchers, work study specialists, psychologists, and designers. This change happened slowly and was driven by new types of challenges that were the result of the introduction of new technologies at the interface between human operators and automated systems. With the massive deployment of computerized systems into process control that began in the late 1970s and continued to transform all human/machine interaction settings during the 1980s and 90s, Human Factors Engineering shifted focus from physical ergonomics towards psychological human factors and began to address cognitive processes behind reasoning, explanations and expectations. Immediately following the accident, the study into the aftermath of Three Mile Island marked the birth of a new interdisciplinary field in which design converges with research and engineering: Cognitive Systems Engineering (Hollnagel and Woods, 1983). The breakdown of information display and communication in the control room at Three Mile Island alerted designers of information displays and control interfaces to the importance of supporting operators in coping with complexity. This insight was based on studying how operators make sense of events and how they respond to anomalies amidst uncertainty when they monitor and control systems (Rasmussen, 1979; Rasmussen and Rouse 1981; Klein, Orasanu and Calderwood, 1993; Hutchins, 1995a, 1995b; Klein and Zsambok, 1996; Vincente, 1999; Woods and Hollnagel, 2006).

Human-computer Interaction (HCI) as a focus in Computer Science and Engineering, Psychology, and Informatics marks one of the most important movements over the past two decades towards a multidisciplinary science for interactions supported by computerized systems (Card, Moran and Newell, 1983; Carroll, et. al., 2003). In parallel, the formation of Usability Engineering (Rosson and Carroll, 2002) is a response to the need for assessing the effectiveness of interactions during a concurrent, iterative design process. Concurrent design practice (Roesler et. al., 2005) illustrates how a diverse set of engineering, research, and design skills is a central requirement in the development of interaction technologies. One central thread in current advances in HCI is the role of visualization in the communication of complex information, in particular in all instances where information is driven by dynamic change over time. But visual form alone is not the only reason why HCI developers are beginning to collaborate closer with designers—designers contribute new design techniques to the development of software and complex technology that provide early evaluation and feedback for design concepts before realization commitment is made (Buxton, 2007).

While Visual Communication and Industrial Designers are experts in establishing visual systems for non-verbal communication, the emerging field of Interaction Design trains visual designers in developing interaction sequences that are useful, usable and understandable. Interaction Designers draw insight from participatory design techniques such as contextual field studies, rapid prototyping and feedback from concept evaluation in a concurrent and iterative design process. Another area of design expertise is the inherently human-centered perspective in the design of technology systems to elicit understanding of how practitioners and users make sense of new interfaces and interactions. To develop interactive systems in a participatory design approach with prospective users, designers apply rapid prototyping and early mock-ups combined with scenario-based design techniques such as storyboards and video prototypes based on storytelling to elicit feedback about designs in progress. Insight from design field studies, but also empirical findings from experimental studies in HCI and Cognitive Science (Hutchins, 1995b) provide the linkage between research, design and technology. The design of interactions in control rooms is driven by the dynamics of work, comprised by the expertise of operators and advisers, operations procedures and the technology that constraints represented processes and representation media.

Nuclear power station control rooms are a perfect example for a high-stakes environment where design requires an integration of these various design, research and engineering fields—a context where design challenge is characterized by high workload and high cognitive demands. Trained operators apply expert knowledge to interpret current situations, calibrate plans for actions and assess possible future consequences of interventions. From a research perspective, expert domains such as control rooms are well studied and documented, since they present highly structured environments where design and operations failures have severe consequences. Training materials, operations procedures and secondary research are available to guide better design practices and avoid error.

The following section on the sequence of events of the Three Mile Island nuclear accident illustrates the complexity of a high-stakes environment work setting. The accident marks one of the well-documented design failures of interactions between human operators and a complex technological system—a case that illustrates the dynamics of operations during anomaly response. Communication failures that emerged in the course of the accident cannot be attributed to a single design error, nor could the overall failure have been prevented by a single design intervention. The Three Mile Island accident was a systems breakdown that emerged from the alignment of several criti-

cal factors that created an extremely brittle situation in which it was easy to make mistakes. As such, the design problems of the control room illustrate the need for a systematic and interdisciplinary design approach to respond to the multifaceted needs for a better interaction design of process monitoring and useful, understandable and usable information representation.

At the advanced level of information interpretation—such as is captured in the following sequence of events—it is crucial for designers to understand the relationship between information representation, the behavior of the system represented and interpretation strategies that operators apply to find explanations that lead them to coordinate interventions.

THE SEQUENCE OF EVENTS

What went wrong in the control room of Three Mile Island? Many accounts of the accident at Three Mile Island-2 provide oversimplified descriptions of the accident. This is endemic to analyses of complex systems failure (Feltovich, Spiro and Coulson, 1997). When analysts fail to understand the intricacies of a system, they may adopt the point of view that nothing is wrong with the system and resort to 'human error' as a cause for what went wrong with a system that performed as designed. To understand the nature of design error in the TMI-2 control room, it is important to understand what caused the complicated situation and how cascading events had rendered information display in the control room insufficient, which in turn led to misinterpretations and confusion about what was happening. The operators in the control room were forced to make decisions in the face of uncertainty.

In a nuclear power plant, a pressurized water system transports heat from the reactor to electric generators (figure 4, numbers in the following refer to this diagram). In the course of this transport, hot, pressurized water from a contained primary cooling system streams through the reactor core (1) and heats water in a (non-nuclear) secondary system that is coupled via steam generators (3) with the primary system. Steam from the steam generators runs a steam turbine (4) and in turn runs the electricity generators (5) for electrical power generation. After giving off heat in the turbines, the steam is condensed (6), cooled down in the cooling towers (7) of the plant and fed back (8) to the steam generators, where it is heated and converted into steam by the heat from the primary cycle.

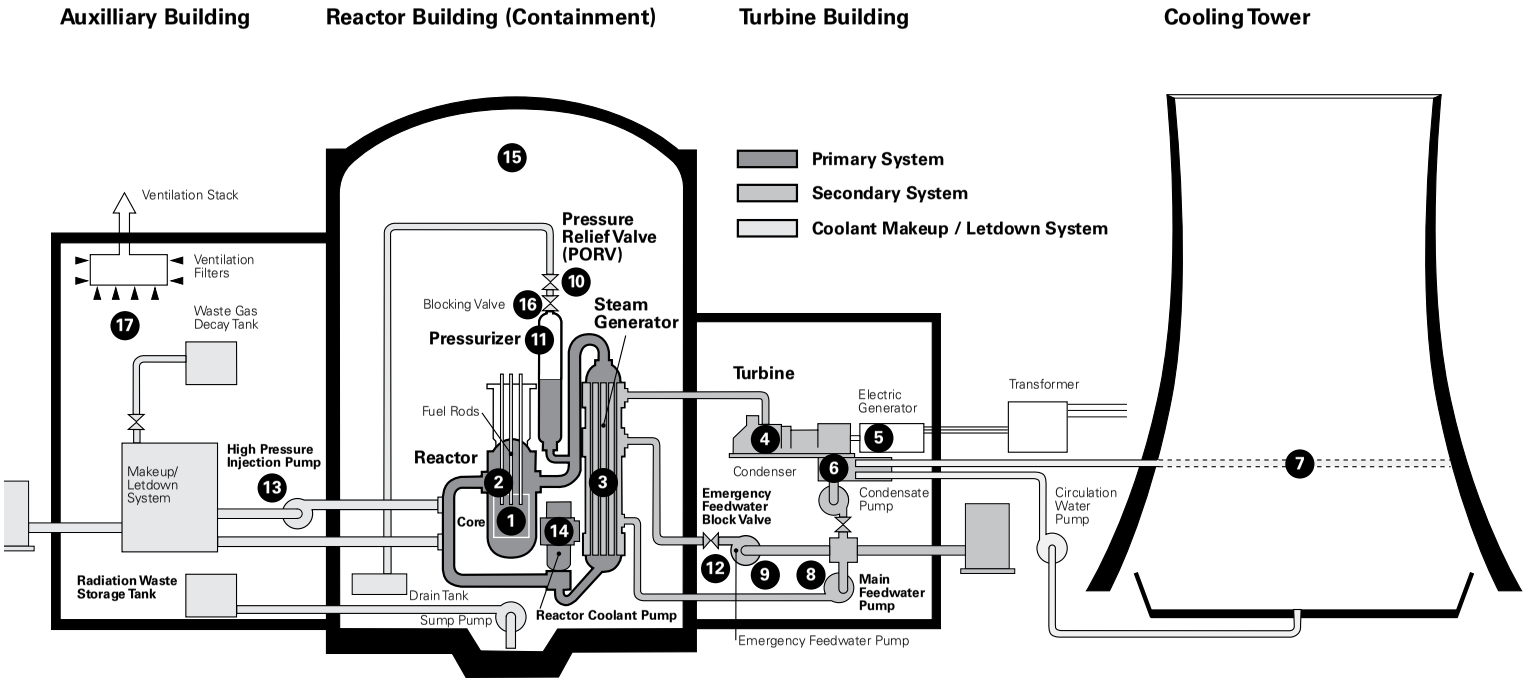


Figure 2
Schematic of the Three Mile Island-2 nuclear power plant (adapted and reprinted from IEEE Spectrum 16, A special report (November 1979) - (© 1979 IEEE).

The Three Mile Island accident began five hours into the graveyard shift. At 4:00 a.m. on Wednesday morning, March 28, 1979, the two main feedwater pumps (8) in the secondary cooling system for the reactor had been shut down by plant automation. Water flow, into the air system of the plant during maintenance activities in the hours before the accident, might have led the automatic tripping of the pumps. With the feedwater pumps shut off, the turbine (4) was tripped by the plant safety automation to prevent it from boiling dry, as new feedwater couldn't be added into the steam generators to produce steam for the turbine. Emergency feedwater pumps (9) were started automatically. With the loss of the turbine, heat transfer from the primary to secondary system was stopped. This led to an increase of coolant temperature, volume, and pressure in the primary cooling system. A pilot operated pressure release valve (PORV) (10) opened automatically to release pressure. Safety systems automatically initiated a scram (automatic shut-down) of the reactor (2). All this happened within 8 seconds from the main feedwater pump trip.

As a result of the reactor shutdown, temperature and pressure of primary system coolant fell and safe pressure margins were regained. The opened pressure relief valve was designed to automatically close once safe margins were restored—but it did not. Indications in the control room, however, led operators to assume that the valve had closed. Unknown to the operators, the valve status light in the control room (figure 5) indicated not the physical valve state, but the flow of electricity that was applied to open the valve. As soon as the electric current to the valve was switched, the control light on the control room console went off. The valve, which normally would close automatically, was physically stuck in the full-open position. This pressure relief valve would play a critical role in the events that followed.

The water in the primary cooling system was normally contained under high-pressure and very high heat. The high pressure prevented the water from boiling at very high temperatures. With the release valve still open, the primary cooling system lost pressure. Unknown to the operators at this point, the secondary system was not operating as expected, as the emergency feedwater pumps (9) were not able to draw water due to shut downstream valves (12). As a consequence, less heat was extracted from the high temperature primary system, and pressure fell as a result of the stuck-open pressure relief valve. Plant safety features automatically initiated high-pressure injection pumps (13) into the primary system to restore high pressure. During this high-pressure injection, the operators monitored the coolant level in a pressurizer vessel (11). They were instructed during training to avoid a situation in which the pressurizer would

completely fill up with coolant, and a rise in pressurizer water level was expected during high-pressure injection. The stuck-open pressure relief (10) valve was located at the top of the pressurizer vessel (11), and the pressurizer steam bubble that under normal circumstances would prevent the pressurizer from filling with coolant had escaped though the open valve. The pressurizer vessel was filling fast and to avoid the pressurizer from going solid (filling up completely with coolant), the operators bypassed the plant emergency system and throttled the high-pressure injection flow. As a result of the throttling of the pumps, pressure in the primary cooling system could not be restored and continued to fall. Although the reactor was tripped, it still produced decay heat, which increased the temperature of the coolant. Now at low pressure and high temperature, steam bubbles began to form in the primary system around the reactor. Steam has a lower heat conductivity than liquid water, which reduced the cooling of the reactor. The expanding steam bubble at the top of the reactor vessel also led to an expansion of the coolant volume—forcing liquid coolant to rise into the pressurizer vessel.

The operators were confused by a high coolant level reading in the pressurizer vessel and a simultaneously low pressure in the primary cooling system. They were trained to take the water level in the pressurizer vessel as an indicator of coolant availability for core coverage—their first priority in the monitoring of the primary cooling, as the core had to be covered with coolant to extract decay heat. Under current conditions, where steam bubbles were forming in a low-pressure, high-temperature coolant—a condition referred to as saturation, pressurizer coolant level, was not an appropriate indicator for coolant inventory. Routine training did not prepare the operators for this assessment, nor for operating the primary cooling system under saturation. They did not have a clear indication that steam was forming in the primary cooling system. They were not aware that the core was undergoing exposure. No instrumentation was available to indicate that the core had entered a particularly critical state.

The high pressurizer coolant level and low pressure in the primary cooling system could have been explained by the fact that the pressure relief valve was stuck open and the pressure was escaping through this steam leak—but the extinguished valve status indicator light on the control room console wrongfully indicated that the pressure relief valve was closed.

Hundreds of alarms sounded in the control room during these first three minutes of the accident *(figure 3)*. In the meantime, the operators realized that the secondary system backup feedwater pumps (9), which were activated by the safety automation after the initial feedwater pump failure, were not able to draw feedwater

because of closed downstream valves (12), which were accidentally left shut after a maintenance activity the day before. The lack of flow was detected 5 minutes into the accident and restored at approximately 8 minutes, but during this time the secondary system side of the steam generators boiled dry. The heat transfer from primary to secondary system was lost.

For the next hour, operators were consumed with bringing the secondary system back into operation. From previous experiences they knew that the secondary system tended to behave unreliably, while the primary system was considered robust. This understanding led them to focus on the recovery of the secondary system, which was complicated by problems in the condenser system (6).

At this point in the accident, the primary system had been in saturation for 54 minutes. The reactor coolant pumps (14) were pumping a two-phase mixture of steam and liquid coolant. This led to reduced flow rates and massive vibrations of the reactor coolant pumps. The operators determined criteria for shutting down these pumps and the criteria were soon exceeded. The operators were now faced with a conflict in keeping the pumps running to maintain primary coolant circulation or shutting them down to prevent losing the pumps altogether due to possible mechanical failure caused by vibration. The pumps were stopped 74 and 101 minutes into the accident, leaving the reactor coolant system without forced circulation. Attempts to establish natural circulation were unsuccessful due to the two-phase combination of the coolant and the low pressure in the primary system; little heat transfer was provided by the slow circulation.

The picture emerged that a steam bubble had formed in the primary system. Upon expansion, the steam bubble began isolating the core from coolant—this led to core exposure and initiated a partial core meltdown. But the operators were not aware of the beginning meltdown, as there were no sensors in the reactor core that could alert them, and no instruments available on the control boards that would directly indicate the health of the reactor. All temperature sensors in proximity to the core were indicating off-range readings of temperature. The operators tried to get information about the reactor status by integrating several information sources, but this process was both difficult and imprecise.

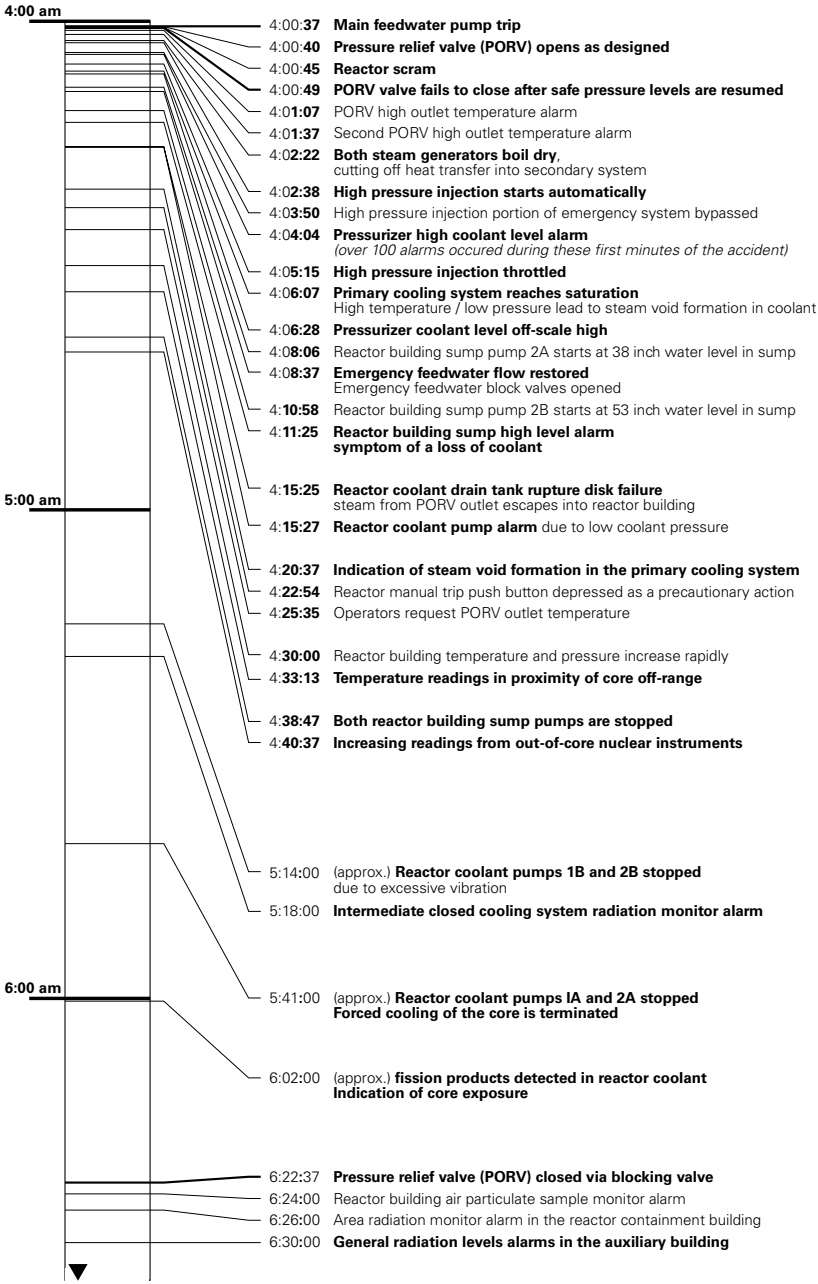


Figure 3
Timeline of a simplified version of the sequence of events during the initial 150 Minutes of the accident. Note the cascading of key events during the initial 5 minutes of the accident (For the detailed sequence of events, see Kemeny (1979), pp. 91-116, available online at http://www.threemileisland.org/downloads//195.pdf).

As a consequence of the beginning reactor meltdown, radiation levels increased in the reactor containment building due to steam that continued to be vented through the stuck-open pressure relief valve while primary coolant water became increasingly contaminated with radioactive products released from the melting core. The first radiation alarms in the containment building (15) were triggered at about 150 minutes into the accident.

The operators had realized, 138 minutes into the accident, that the pilot-operated pressure release valve (10) did not close automatically as designed—it was opened 3 seconds into the accident and was supposed to be closed 9 seconds later. Using a back-up blocking valve (16), they closed the PORV manually. This and another series of interventions allowed them to stabilize pressure and reactivate one of the emergency high-pressure injection pumps (13) to add coolant and recover the core. In the meantime, radiation from the containment building had spread into the adjacent Auxiliary Building (17) and had been distributed across the plant. A site emergency was declared at 6:56 a.m. due to severe radiation levels in the entire plant, followed by the declaration of a general emergency at 7:24 a.m. Over the course of the day, in a series of high-pressure injection flows and various depressurization procedures, the operators were able to re-activate circulation in both the primary and secondary cooling systems, leading to a stabilization of the reactor cooling system at 7:50 p.m.—15 hours and 49 minutes after the events that initiated the accident.

Even then the problems were not yet over, and in fact, worsened over the following days due to the formation of a hydrogen bubble in the reactor pressure vessel. The hydrogen bubble was the result of a chemical reaction between the exposed core and coolant. Experts were concerned that the bubble could lead to an explosion in the reactor vessel, posing the threat of a massive release of radioactivity into the environment and requiring the evacuation of 100,000 people in the area 20 miles downwind from the plant. This worst-case scenario fortunately did not become a reality, as there was no oxygen in the reactor present to make an explosion possible. During the assessment of the dangers of the hydrogen bubble, experts had found themselves confronted with a new situation, as very little was known about the science of a reactor after extensive core exposure.

It would take almost a month before the reactor was stabilized. Analysis of the reactor core nine years later in 1988 revealed that one half of the core had melted during the early stages of the accident. Many experts were surprised that the reactor vessel had not been breached by the massive meltdown. Today, the reactor remains

nonoperational. The damaged fuel was removed in 1993, and the reactor is today in monitored cold storage operation. TMR-1, the other reactor of the plant, which had been shut down for refueling shortly before the accident, remained closed after the accident, but resumed operation in 1986.

THE TMI-2 CONTROL ROOM AS AN OPERATIONS-CRITICAL SYSTEM

Control rooms impose unique design challenges: information that is required to steer processes in the event of a disturbance needs to be presented during the situation when it is needed and in a format that is useful.

The design of a nuclear power plant makes direct observation of events impossible. Buildings, pipes and barriers obscure the key processes, that drive critical conditions. Configurations are large in size and located at a distance. The reactor core is contained in a 36 foot tall reactor pressure vessel made of 9 inch thick steel. It contains 100 tons of nuclear fuel. Two steam generators, each 72 foot tall and the reactor cooling system, comprised of room size pumps, man hole diameter pipes, storage tanks and blocking valves, are controlled remotely. To monitor information from a complex system of components that are spatially distributed but functionally related, thousands of sensors are located across the plant. At Three Mile Island, 2,400 displays represented the read-outs of these sensors.

The control room at Three Mile Island-2 was the product of a bottom-up engineering approach. Because the designers of the plant had assumed that a number of particularly critical operational states were virtually impossible, many displays that might be required to cope with such "impossible" states were not included in the control room design. Instead of integrated displays that indicated the overall safety of the plant, the boards in the control room provided operators with thousands of displays and alarms. Each showed detailed states of separate components. Then the so-called 'impossible' happened.

In an attempt to patch together a larger picture of what was occurring, operators had to compare several displays, sometimes placed in random locations across the 900-square-feet of control boards—all this in the presence of hundreds of alarms being activated one after another first by a root cause, leading to follow-up alarms in the rapid pace of subsequent events that by then were already irreversible. Hundreds of alarms sounded within the initial minutes of the accident. The design of the control room made no consideration for how fundamentally different the operational demands would be during a non-typical situation such as this emergency.

During the formation of steam bubbles in the primary cooling system (an event that transformed operations into a status that was outside operations procedures), the control room performed as designed. Post-accident assessments of Three Mile Island from an engineering perspective state that all information was available in the control room to initiate the appropriate responses (Kemeny, 1979). Interviews with control room operators, however, tell a different story: Bombarded with hundreds of alarms and flashing lights during the first minutes of the accident, the operators were unable to distinguish important information from irrelevant data. Comment of one of the operators on duty during the accident in the early morning hours of March 28, 1979. quoted from the official inquiry following the Three Mile Island accident (Kemeny, 1979): *"I would have liked to have thrown away the alarm panel. It wasn't giving us any useful information."*

Confusion arose because detail of singular data had been favored over establishing relationships between data in context. During advanced stages of the accident, available information was conflicting and desired information was not available. The designers of the control room had left it to the operators to identify what was important and to ignore what was irrelevant. In the course of several events that required responses, data from thousands of sensors in the plant was displayed on the control boards, but all this data was presented with no particular hierarchy, meaningful spatial association or context showing its relationship to other data. A computer printer used to log all data read-outs was running hours behind, and the operators had to dump its memory several times to print out reports that they needed (Rogovin, 1980).

Key processes during the development of the accident could not be detected because temporal plots of data trends were not available and a big picture explanation of what was going on depended on the manual correlation of related data streams that were displayed on spatially separated control boards. At the core of these observations lies a dynamic visualization problem that resulted in an ineffective, opaque representation of events at a distance.

What had happened at Three Mile Island was a dire combination of disturbances and failures that the designers of the plant thought was impossible. Insufficient representation of the relationship between these failures, and trends of change resulting from the failures, obscured the severity of the situation.

The accident investigations developed alternative scenarios of the accident that pointed out that the development from disturbance into a serious accident could have been prevented if the operators had been able to detect the wrongfully open release valve (Kemeny, 1979). But this view reflects a typical hindsight bias: In the aftermath of a

mistake, one can point to the causes of the error, as all consequences are known and the complete function of the system that has failed has been determined. What made the accident at Three Mile Island so difficult to control as it developed was that the operators were confronted with unclear, incomplete and conflicting information that was insufficient to alert them that they were dealing with a potentially dangerous situation well outside the safe boundary conditions for operations of the plant.

Other significantly complicating factors were the pacing and cascading of events, alarm escalation and the inadequate mapping between physical and functional relationships in the display of data. The combination of these factors led the operators to the formulation of a wrong mental model during anomaly response; this illustrates a deep conflict in the design philosophy of the control room: The control room was designed for normal operations—this rendered it insufficient for the display of operations-critical information during an anomaly when this information was needed the most.

ALARM ESCALATION RESULTED IN DATA OVERLOAD

Within seconds after the initial events of the accident, the control board lit up like a christmas tree. Important information was masked by the sheer quantity of hundreds of alarms that were triggered one after another (Woods, 1995). It was hard to distinguish important information from less relevant data, as critical pieces of information (indicator lamps for the closed valves downstream from the emergency feedwater pumps in the secondary system, for example) were masked by data from follow-up alarms. The operators were consumed by tracking side effects and follow-up complications—this sidetracked them from diagnosing the bigger situation they were in.

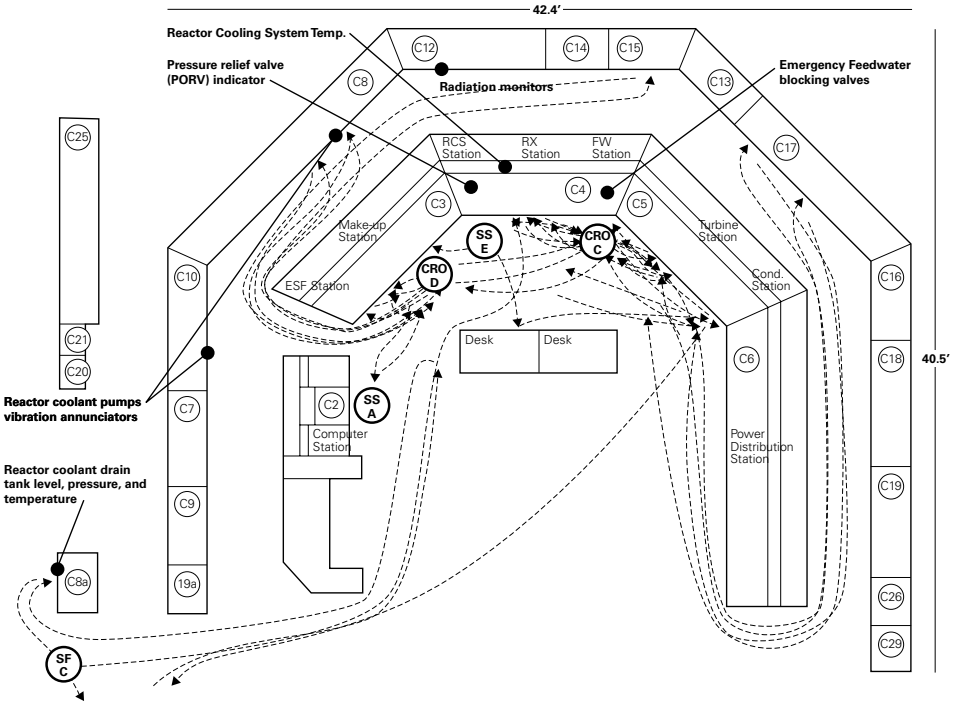
A number of the many alarms that sounded were known to be nuisance alarms, indications of unusual readings that were not operations critical. As a result, other critical alarms were dismissed. One example of this is the reading of the downstream temperature of the stuck-open pressure relief valve. The post accident reports point out that the operators should have realized that the valve was open by noticing an increased temperature reading downstream from the valve. But the operators knew that the valve was leaky, and they were used to high temperature readings downstream from the valve. Since the high temperature reading was discounted as a nuisance alarm, it masked the open state of the very critical faulty pressure relief valve.

INADEQUATE MAPPING BETWEEN PHYSICAL AND FUNCTIONAL RELATIONSHIPS IN THE DISPLAY OF DATA LED TO CONFUSION

The control room of the Three Mile Island nuclear power plant provided the interface for both information display and control interactions with the critical condition of the reactor. The control room provided the visual work environment for the team of operators: two control room operators, a shift supervisor and a shift foreman, who together needed to gather a shared understanding of the situation that was unfolding. This understanding was crucial to formulate plans and implement actions in order to bring the unstable state of the reactor and cooling system back to stable operations. The layout of the control room and arrangement of control panels required the operators to relocate their positions many times during the development of the accident (figure 4). Several instruments that required correlation with other instruments could not be monitored simultaneously, as they were located too far apart or were obstructed from the view of the operators by the control console (Malone, 1980). (For more photographs of the TMI-2 control room, including detail views of displays and controls, see the Human Factors review in Rogovin, 1980, Volume 2, Part 2, 573-612, available online as PDF at http://www.threemileisland.org/downloads//202.pdf.)

Several times during critical phases of the accident, operators came to wrong conclusions based on conflicting information displays. Comparison of separate displays was not only complicated by the location of the instruments, but also by the escalating temporal sequence in which changes of data occurred. As an example, backup information cues, that indicated the pressure relief valve indicator light, was a wrongful indication of the valve status and it came from eight different sources, seven of which were indicated within the first and most dense 4 minutes and 30 seconds as the accident unfolded. Comparison between instruments was complicated by the random arrangement of displays and controls. The eight cues were scattered across the 900-square-feet of control panels, requiring the operators to correlate separate displays that were located on distant panels amidst hectic activities in the control room. One of the eight cues was a level indicator for the reactor coolant drain tank overflow. The indicator light was located at the rear of the control panel *(figure 4, C8a)*. The arrangement of displays and controls could have been greatly improved by a better mapping between the displays and the physical location and functional relationship of their corresponding sensors in the plant.

Three Mile Island Unit-2 Control Room
Panel arrangement and operator re-location during the initial 150 minutes of the accident



Control Room Features (approx. numbers)

Panel space	**900** sq ft
Controls	**1200**
Displays	**2400**
Integrated controls/displays	**1500**
Annunciators	**800**
Maximum viewing distance	**48** ft
Maximum walking distance	**50** ft

Acronyms:

ESF - Engineered Safety Feature
FW - Feedwater
RCS - Reactor Coolant System

CRO - Control Room Operator
SS - Shift Supervisor
SF - Shift Foreman

Panel Layout

Inner Consoles

C2 Computer console
C3 Auxiliary systems console
C4 Plant control console
C5 Turbine control console
C6 Electric control console

Back Panels (located outside of main control area)

C8a Reactor coolant drain tank panel
C20 Nuclear instrumentation cabinet No. 1
C21 Nuclear instrumentation cabinet No. 2
C25 HV & AC panel

Vertical Panels

C7 Fire detection panel
C8 Coolant systems monitoring panel
C9 Push-pull control panel
C10 Plant equipment temperature recording panel
C12 Radiation monitoring panel
C13 Engineered safety features panel
C14 Control rod drive panel
C15 Containment isolation panel
C16 Turbine supervisory panel
C17 Turbine auxiliary monitoring panel
C18 Staion electric auxiliary monitoring panel
C19 Vital power panel
C19A 500 KV control panel
C26 Diesel generator No. 1 panel
C29 Diesel generator No. 2 panel

Figure 4
Layout of the Three Mile Island-2 control room. Panel arrangement and operator re-location during the initial 150 minutes of the accident.

A better visual organization of the information displays with an integrated approach, showing relationships rather than isolated instances, would have greatly supported the operators' assessment of the most critical events that were going on unnoticed: reactor core exposure and subsequent partial meltdown. But the design relied on the operators to accomplish this integration. In the presence of everything going on, it was unlikely that the operators would have been able to construct an accurate 'big picture' of what was going wrong. The operators were asked to do the impossible.

THE WRONG MENTAL MODEL

Instead of attributing the overheating of the primary system to the fact that too little coolant was present (because steam escaped through the stuck-open pressure relief valve), the operators thought an abundance of coolant was present as a result of adding too much coolant during high pressure injection flow. This wrong mental model led them to underestimate the risk of reactor core exposure. The emergence of this wrong mental model can be explained with the information available at the time in the control room and instructions provided by operations procedures: The most unusual condition during critical phases of the accident was the indication of a high level of coolant (in the pressurizer vessel) and low coolant pressure—this was a set of indications that conflicted with the operations procedures that identified a loss of coolant accident (the type of accident that could lead to core exposure) as a situation during which operators would experience both low level of coolant and low pressure in the reactor cooling system. Instead, at Three Mile Island the high coolant levels at low temperature readings were due to a steam bubble formation in the low-pressure coolant. Under this condition, following operating procedures that used pressurizer level as an indicator for coolant inventory (in other words, whether the core was sufficiently cooled) would not apply (Malone, 1980).

Using the pressurizer level as an indicator of sufficient coolant inventory at the core led the operators to misinterpret the off-range high temperature readings of instruments close to the core. The instrument had started to display off range readings 6 minutes into the accident, but the operators thought that these temperature readings weren't off range, but that the analog instruments had stabilized close to the max reading. They thought core exposure wasn't possible due to the high level coolant indications, but they didn't know that steam formation had taken place, because an instrument that would integrate pressure and temperature of the coolant was not available in the control room.

There were indirect indications that saturation had taken place—these included severe vibrations read-outs from the reactor coolant pumps approximately one hour into the accident. Again, data overload made the detection of this excessive vibration difficult, because the instruments were buried in the large number of other displays and active alarms.

KEY SAFETY INFORMATION WASN'T AVAILABLE

However much the bias of hindsight would argue that all the information to make the right decisions had been available (it had just been difficult to see it on all the vast arrays of instruments and alarms that were constantly sounding), there was one central piece of information that was missing and that the operators could only infer from a number of proximity sensor read outs. The question was 'do we have coolant in the core'? This piece of information was central to the events going on because in light of the wrong mental model, the operators thought that enough coolant inventory would be available to extract heat from the core. But, unknown to them, a steam bubble had formed, exposed the core and pushed coolant into the pressurizer vessel.

Direct sensors at the core, coupled with alarms, would have alerted the operators that the core was exposed. The extreme conditions in the core made this type of sensing expensive, and the plant design assumed that the information of core status could be deduced from sensor read-outs in proximity to the core. As the accident progressed and the operators began to realize that steam formation had been under way, they began to suspect the presence of a steam bubble, but with the data they had available, they could not determine where the steam bubble was.

PLANT OPERATIONS LEFT SAFE BOUNDARY CONDITIONS

The control room was not designed to alert the operators that they had left the safe boundaries of operations. There was no indication in the instrumentation panels that plant operations were in the middle of a dangerous anomaly. During critical sequences in the accident development, the operators didn't realize that they needed the help of experts to cope with the situation they were in. They didn't know that operations procedures were no longer enough to bring the plant back into safety. Distorted information in the control room prevented complete awareness of the situation they were in. Communication lines between the control room and outside world were limited. Several times during the accident, reporters dialed directly into the control room.

The operators were not trained in operating the plant under the conditions they were facing. Analysis after the Three Mile Island accident showed that the control room simulator that was used to train control room personnel was not able to simulate the type of loss-of-coolant accident that actually happened (Kemeny, 1979). In fact, the computing power of the simulator was too limited to simulate the complex sequence of cascading events that affected the thermodynamic system during the accident. Nobody thought an accident of this kind was possible.

CONCLUSION

The larger portion about what we know today about human-computer interaction in high-stakes domains is the result of studies of design efforts gone wrong. Three Mile Island is one example of a number of accidents that have been thoroughly studied and documented. Understanding design error in these cases has produced new knowledge that positions us to better deal with design error in the future.

Design work in technology-intense, high-stakes domains requires the understanding of the reasoning strategies that human operators apply while working with computerized systems that represent tasks at a distance. This entails how the represented system adapts to change during interventions and how its constraints drive interactions—leading the resulting joint cognitive system response to change (Hollnagel and Woods, 2005). The complexity of the design challenges for supporting interactions between operators and control systems reflects the complexity of any controlled system itself: events, pacing and multiple perspectives of operators and automated systems during interventions are all subject to dynamic change and changes

are coupled across several layers of functional relationships in the affected system. Initial change can affect multiple relationships in a system in parallel. Interruptions can lead to cascading events and changes—both fast and slow—that are difficult to control (Woods and Hollnagel, 2006). On the other hand, very slow changes that affect the overall system significantly might be difficult to detect and/or control. They might form unnoticed, if the data change over time is progressing too slowly to be detected, or if changes that unfold are located outside the frame of view defined by the perspectives of operators and automation.

Control rooms align representation properties with systems properties and confront both with the challenges of mediated control at a distance. Control room information displays represent the status and trends of the remote processes that are being monitored. Operations in control rooms are the product of both human reasoning and machine support. To design effective representations of operations, we need to understand how operators form explanations based on the display of data. To do this, we need to fully consider the situation into which data is presented.

Well-documented cases of design error illustrate the role of context—environments, individual views, shared understanding and critical situations during failure; they form a knowledge base of control systems breakdown. Examples are the reports of the National Transport Safety Board (www.ntsb.gov) on aviation accidents or the case of the Therac-25 accidents with a malfunctioning radiotherapy device that led to several fatalities during cancer treatment resulting from radiation overdoses due to software failure (Leveson and Turner, 1993). Murray and Cox (1989) and Mindell (2008) capture several instances of information display and automation failures in their history of mission control during the Apollo space missions. As a short list of interaction design failures in high-stakes environments that might particularly appeal to designers, these case studies illustrate the limits of technology-centric development and demonstrate the need for bringing more design expertise into human-computer interaction development teams to elicit the views of operators on the envisioned systems before the systems are realized—providing opportunities for developing the right systems the right way—which will result in artifacts and representation design systems that support the cognitive work of practitioners and help them cope with complexity.

Lessons from design error mark turning points in the design of automated control systems and have led to revised views about the roles of designers, human operators and computerized systems in process control. The popular attribution of 'human error' oftentimes is nothing more than an oversimplified account of cause and responsibilities—it provides little insight to improve the designs that lured operators into making wrong assessments. The conception of 'human error' contradicts the reality of the complexities and stakes of operations where expert practitioners work with technology-intense systems. When encountering error, these operators try the best they can to get the situation back under control. They are at the sharp end of the process, and they will be exposed to its consequences first and held responsible for operations gone wrong later (Woods et. al., 1994).

Designers on the other hand, are stakeholders from a distance. Any designed artifact that supports practitioners in making decisions is a stand-in for the designers of the system. The designers are responsible for their systems, even if they are not on site to guide operators through bottlenecks and challenges. Problems with the designed system, however, can be expected. Things that can go wrong will go wrong, and any artificial system can eventually be involved in a situation that is outside the terrain for which it was designed.

Therefore we must both anticipate error and design for error. Design for resilience supports operators in coping with errors by moving the likelihood of errors into protected territories in the design rather that making critical aspects of the design prone to error. This is especially important where outcomes of errors might be difficult to steer (Hollnagel and Woods, 2006). A good control room interface should provide operators with a 'big picture' of the situation the system is in. It should provide information of current status and trends, and it should indicate how the current status relates to expected outcomes, stable operations and the safe boundaries of operations. A good control room interface should not allow operators to steer a process into a catastrophe—in the worst case it might allow them to severely damage the plant. From this perspective, the Three Mile Island control room performed as designed. It reflected design and engineering knowledge at the time and its failure has laid the foundation for new contributions to knowledge in the aftermath of the Three Mile Island accident in 1979.

Newly designed systems will no doubt introduce new errors. Response strategies to cope with these errors may be as novel as the design that produced them, therefore, strategies for coping with error have to be part of the design.

REFERENCES

Adams, D. 1992. *Mostly Harmless*. New York, NY: Harmony Press.

Buxton, W. 2007. *Sketching User Experience: Getting the Design Right and the Right Design*. San Francisco, CA: Morgan Kauffman.

Card, Stuart K., Thomas P. Moran and Allen Newell. 1983. *The Psychology of Human–Computer Interaction*. Erlbaum, Hillsdale, NJ: Lawrence Erlbaum and Associates.

Carroll, J. M. 2001. *HCI models, theories, and frameworks: toward a multidisciplinary science*. San Francisco, CA: Morgan Kaufman.

Feltovich, P.J., R.L. Coulson and R.J. Spiro. 2001. Learners' (mis)understanding of important and difficult concepts: A challenge to smart machines in education. In Forbus, K.D. and P.J. Feltovich, editors. *Smart machines in education*. Menlo Park, CA: AAAI/MIT Press.

Hollnagel, E. and D.D. Woods. 1983. Cognitive Systems Engineering: New wine in new bottles. *International Journal of Man-Machine Studies*, 18, 583-600.

Hollnagel, E. and D.D. Woods. 2005. *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. London, UK: Taylor & Francis.

Hollnagel, E., D.D. Woods and N. Leveson. 2006. *Resilience Engineering: Concepts and precepts.* Aldershot, UK: Ashgate.

Hutchins, E. 1995a. *Cognition in the wild*. Cambridge, MA: MIT Press.

Hutchins, E. 1995b. How a cockpit remembers its speeds. *Cognitive Science, 19,* 265-288.

Kemeny, J.G. 1979. *Report of the President's Commission on the Accident at Three Mile Island*. New York, NY: Pergamon Press. (http://www.threemileisland.org/downloads//195.pdf Volume 1, additional volumes at http://www.threemileisland.org/resource_center/index.php).

Klein, G., J. Orasanu and R. Calderwood. 1993). *Decision Making in Action: Models and Methods.* Norwood, NJ: Ablex.

Klein, G. and C. Zsambok, editors. 1996. *Naturalistic Decision Making.* Mahwah, NJ: Lawrence Erlbaum and Associates.

Leveson, N.G. and C.S. Turner. 1993. An investigation of the Therac-25 accidents. *Computer*, July, 18-41.

Malone, T.B. and U.S. Nuclear Regulatory Commission, Special Inquiry Group and Essex Corporation, et. al. 1980. *Human Factors Evaluation of Control Room Design and Operator Performance at Three Mile Island-2: Final Report.* Washington, DC: The Group.

Mindell, D.A. 2008. *Digital Apollo: Human and machine in space flight*. Cambridge, MA: MIT Press.

Murray, C. and C.B. Cox. 1989. *Apollo, The Race to the Moon.* New York, NY: Simon & Schuster.

Rasmussen, J. 1979. *On the Structure of Knowledge: A Morphology of Mental Models in a Man-machine System Context*. Risø-M-2192. Roskilde, DK: Electronics Department, Risø National Laboratory.

Rasmussen, J. and W.B. Rouse. 1981. *Human detection and diagnosis of system failures*. New York, NY: Plenum.

Roesler, A. and D.D. Woods. 2005. *Inventing the future of cognitive work: Navigating the northwest passage, Proceedings of the 6th international conference of the European Academy of Design*. University of the Arts, Bremen, Germany, March 29–31 2005.

Rogovin, M., U.S. Nuclear Regulatory Commission. Special Inquiry Group. 1980. *Three Mile Island: A Report to the Commissioners and to the Public*. Washington, DC: U.S. Nuclear Regulatory Commission. Special Inquiry Group. (http://www.threemileisland.org/downloads//354.pdf and http://www.threemileisland.org/downloads//202.pdf Volumes 1 and 2, additional volumes at http://www.threemileisland.org/resource_center/index.php)

Rosson, M.B. and J.M. Carroll. 2002. *Usability Engineering: Scenario-based development of human-computer interaction*. San Francisco, CA: Morgan Kauffmann.

Vicente, K.J. 1999a. *Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-Based Work*. Mahwah, NJ: Lawrence Erlbaum and Associates.

Walker, S. 2004. *Three Mile Island: A Nuclear Crisis in Historical Perspective*. Berkeley, CA: University of California Press.

Woods, D.D., L. Johannesen, R.I. Cook and N. Sarter. 1994. *Behind human error: Cognitive systems, computers and hindsight*. Dayton OH: Crew Systems Ergonomic Information and Analysis Center, WPAFB. (http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA492127&Location=U2&doc=GetTRDoc.pdf)

Woods, D.D. 1995b. The alarm problem and directed attention in dynamic fault management. *Ergonomics, 38*, 2371-2393.

Woods, D.D. and E. Hollnagel. 2006. *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. London, UK: Taylor & Francis.