There Is More to Monitoring a Nuclear Power Plant than Meets the Eye

Randall J. Mumaw, Boeing Commercial Airplane Group, Seattle, Washington, **Emilie M. Roth**, Roth Cognitive Engineering, Brookline, Massachusetts, **Kim J. Vicente**, University of Toronto, Toronto, Canada, and **Catherine M. Burns**, University of Waterloo, Waterloo, Canada

A fundamental challenge in studying cognitive systems in context is how to move from the specific work setting studied to a more general understanding of distributed cognitive work and how to support it. We present a series of cognitive field studies that illustrate one response to this challenge. Our focus was on how nuclear power plant (NPP) operators monitor plant state during normal operating conditions. We studied operators at two NPPs with different control room interfaces. We identified strong consistencies with respect to factors that made monitoring difficult and the strategies that operators have developed to facilitate monitoring. We found that what makes monitoring difficult is not the need to identify subtle abnormal indications against a quiescent background, but rather the need to identify and pursue relevant findings against a noisy background. Operators devised proactive strategies to make important information more salient or reduce meaningless change, create new information, and off-load some cognitive processing onto the interface. These findings emphasize the active problem-solving nature of monitoring, and highlight the use of strategies for knowledge-driven monitoring and the proactive adaptation of the interface to support monitoring. Potential applications of this research include control room design for process control and alarm systems and user interfaces for complex systems.

INTRODUCTION

If I ever took an alarm summary and didn't find anything there, I wouldn't believe it. I would think that there was something wrong with the computer.
　　　Anonymous nuclear power plant operator

One of the fundamental challenges in studying cognitive systems in context is determining how to generalize beyond the specific case studied. How does one progress beyond the details of a single case to derive a more general understanding of distributed cognitive work and how to support it? The ability to generalize beyond particular settings and technologies to identify the deeper factors at work is an important prerequisite for advancing the field of cognitive engineering (Woods, 1995b). This paper attempts to contribute to the understanding of this issue by presenting one approach to the problem of how to generalize from field data.

Our specific case was a set of field studies that examined how nuclear power plant (NPP) operators monitor plant state during normal operating conditions. We describe the rationale for choosing to use field observations to shed light on this issue and the process we used to gain insights that would generalize beyond the specific individuals we observed, the specific

plants we sampled, and even the specific industry. Our goals were to address the specific need to understand NPP operator monitoring during normal operations as well as advance the general knowledge of the factors that can make monitoring difficult and the types of strategies that expert practitioners have developed in response to those demands.

Historically, research in the nuclear power industry has focused on analyzing operator performance in abnormal or emergency situations (e.g., Roth, Mumaw, & Lewis, 1994; Vicente, Moray et al., 1996; Woods, Wise, & Hanes, 1982) – a focus that is easily justified by the enormous impact that such situations can have on human life and the environment. The result of this focus on emergency operations is that investigations of normal operating conditions are relatively rare. Because most abnormalities are preceded by normal operating regimes, an understanding of this process is critical. The operator's role in detecting early signs of a problem and taking appropriate actions to prevent further degradation is a key element in preventing abnormal situations from occurring. Moreover, an understanding of operator monitoring can provide guidance for improving operator interfaces and support systems for normal operations. We undertook a research program directed at beginning to fill this gap in knowledge.

Why Field Observations?

Our research program began with a review of the literature on the factors that affect human monitoring performance (Mumaw, Roth, Vicente, & Burns, 1995). This review included both the classic literature on monitoring performance (e.g., Moray, 1986) and studies and models of monitoring as part of disturbance management (e.g., Woods, 1995a). These sources yielded theoretical constructs that provided the grounding and interpretive framework for our empirical investigations (e.g., the concept that monitoring is affected by both knowledge-driven and data-driven processes). At the same time, the context of normal NPP operations is sufficiently complex and different from the contexts in which monitoring has traditionally been examined that it was unclear how well these theoretical constructs generalized to monitoring during normal operations.

Monitoring demands in a nuclear power plant are very different from those in the types of monitoring tasks that have been examined in laboratory settings. An NPP has thousands of parameters that operators can potentially monitor. Further, monitoring is not the only activity in which power plant operators are engaged; monitoring activity is interweaved with other ongoing responsibilities of managing day-to-day tasks for generating power. The questions of interest were as follows: How do operators routinely monitor the plant? What factors contribute to monitoring difficulty and vulnerability to missing key indicators of a developing problem? What cognitive strategies have operators developed to facilitate early detection of a problem? These questions are important from the perspective of both understanding the contributors to monitoring error and identifying opportunities to improve monitoring effectiveness through the introduction of new aids or training. These questions could best be addressed by examining operator performance in the work setting.

One primary strength of naturalistic observations is that they support a discovery process (Woods, 1995b). They serve to draw attention to significant phenomena and suggest new ideas, the validity and generality of which can then be evaluated through additional studies. The field studies we performed served this role. They helped us to discover and characterize the cognitive complexities that operators confront during normal operations and the active, problem-solving activities that play a fundamental role in monitoring a complex, dynamic system such as an NPP.

Field observations afford the opportunity to gain a realistic view of the full complexity of the work environment, which is not possible with other approaches. We were able to observe the broad range of operator duties, the distractions that arise, the need to time-share tasks, other demands on operator cognition or types of complexity, and the interactions with other people both inside and outside the control room. Second, field observations allow one to observe the distributed cognitive system (Hutchins, 1995), examining how operators

adapt the work environment in order to off-load cognitive demands and facilitate performance. Finally, field observations provide the opportunity to observe and document actual incidents that reveal the cognitive demands operating in the environment and strategies that have been developed by operators in response to these demands. These illustrative incidents may be examples of operator performance in routine situations, or they may represent a response to a relatively rare occurrence (e.g., equipment malfunction; accident) that arises during the field study.

Before we began this research, we (and our sponsor) had several hypotheses about the demands associated with monitoring during normal operations. One hypothesis was that the primary challenge to monitoring is operator vigilance. After all, large-scale accidents, such as those that occurred at Three Mile Island and Chernobyl, are extremely rare, leading to an often-heard characterization of process control tasks as "99% boredom and 1% sheer terror." Minor upsets occur more frequently but are still relatively rare. An alternative hypothesis was that the primary demands in monitoring relate to selective attention. Given that an NPP control room consists of hundreds of indicators, one might think that the difficulty is in choosing what to attend to and what to ignore – a virtually impossible selective attention task. A third hypothesis was that the major difficulty in monitoring is one of visual perception. Because the control room is so large and consists of so many instruments, the key to monitoring may be to develop the visual acuity and discrimination skills that are required to detect changes and read indicators accurately. As the remainder of this paper shows, none of these hypotheses does justice to the richness of the phenomena we observed. Although vigilance, attention, and visual perception are relevant issues, we found that there is much more to monitoring than meets the eye.

The four field studies are described first with an overview of the approach and the objective of each study. The next section is used to describe the work setting and provide a more detailed methodology for each of the four studies. Because each study is used to fur-

ther develop the results of its predecessor study, we present our findings in a single section that focuses on the research questions instead of the studies – that is, the contributions of each successive study are not described in detail. The section "How Operators Monitor during Normal Operations" lays out findings from three perspectives: sources of information for monitoring, what makes monitoring difficult, and the strategies operators use to facilitate monitoring. The discussion section addresses the implications of our findings, offers a broader framework for monitoring, and revisits the issue of generalizability in cognitive field studies.

Overview of Research Approach: A Series of Cognitive Field Studies

The field study methodology we adopted is similar in approach to ethnographic methods (Jordan & Henderson, 1995; Nardi, 1997) and the European field study tradition (De Keyser, 1990) in that we observed and interviewed operators in the work setting. We use the term *cognitive field studies* to emphasize that our focus was on the cognitive and collaborative processes that contributed to monitoring performance and that our observations and analyses were strongly informed by cognitive theories. We acknowledge that our observations are ultimately guided by (and potentially limited by) the cognitive theories and bodies of knowledge that we use in interpreting what we see and determining what is "interesting."

The objectives of the cognitive field studies were to document established monitoring practices and to uncover operator strategies that facilitate effective monitoring during normal operations. We placed ourselves in the actual work environment to observe and interview operators as opportunities arose. In general, for each data collection session, we sat in on most of the 12-hr shift with a single operator (or pair of operators) to experience the various activities that occur over the course of the day (or night). Interviews were used to obtain general information about operator practice, to follow up on monitoring actions we observed, and to test the representativeness of the monitoring strategies we observed or were told about from other operators.

Note that our aim in performing cognitive field research is to uncover the cognitive demands of the environment and the high-level cognitive strategies that practitioners have developed to cope with those demands. Our aim is to specify not the operators' detailed internal mental representations or mental processes but, rather, the strategies they employ and the resources they exploit to cope with domain demands. In this sense our approach is more closely aligned with the distributed cognition tradition (e.g., Hutchins, 1995) than with the cognitivist approach that attempts to specify and model internal mental representations and processes at a more micro-level of analysis (Vicente, 1999).

A primary limitation of field studies is that it can be difficult to assess the extent to which the results generalize beyond the specific case studied. Our strategy to overcome this limitation was to perform a series of field studies explicitly aimed at testing the limits of generalization of our findings. Data collection was conducted in several separate studies using a bootstrap approach. Study 1 was conducted by two observers, each working independently in a four-unit NPP control room. We began with an unstructured set of on-site observations and interviews to familiarize ourselves with the operator's tasks. This more opportunistic approach resulted in an initial characterization of operator monitoring activities during normal operations, the sources of information for monitoring, what makes monitoring hard, and what strategies operators have developed to facilitate monitoring. The subsequent studies involved additional field interviews (unstructured) and observations designed to test, refine, and expand the initial characterization.

In Studies 2 and 3 we returned to the same plant to determine whether our findings were specific to the operators that we had observed or were representative of operators at that plant. We used two complementary methods to assess the generality of the findings. In Study 2 we circulated a written report of our Study 1 findings to the full set of plant operators to solicit their feedback on the validity and generality of our preliminary findings and interpretations. In Study 3 we collected additional observational data on a new sample of opera-

tors at the same plant. In Study 4 we went into a second plant that had a slightly more technologically advanced control room interface. Here, as in Study 3, we used unstructured observations and interviews to assess whether the findings regarding the cognitive drivers of operator monitoring and the activities operators engage in to facilitate monitoring generalize across plants and interface technology.

By using a progression of field studies, we were able to establish that our findings regarding what made monitoring difficult and the active strategies that operators have developed to facilitate monitoring held across individuals, across time, and across plants that differed in sophistication of automation and user interface technology.

FIELD STUDIES

In this section we describe the four field studies and provide a brief overview of the types of results obtained in each study. The primary aim is to explain how each study motivated and guided the next one. Because Studies 2 through 4 were designed to validate and extend the generality of the earlier studies, they replicate and enrich the results of the prior studies. For ease of exposition, we postpone full presentation and discussion of the results until all four studies are described.

Description of NPP Control Rooms

Our research sponsor provided access to two Canadian NPPs. Studies 1, 2, and 3 were conducted in Plant A. This control room has four control units (each controlling its own reactor). Figure 1 shows the panels on a single unit (from Plant B, which is similar). A single operator runs each unit, although other personnel serve support roles. Each control unit occupies a "corner" of a single, large room that is completely open (i.e., no barriers to visibility). Therefore, the operator on each unit can see the panels and alarms of all other units. This allows operators to monitor activity on other units and maintain an overall awareness of plant activity.

Each unit consists of stand-up control panels, an operator desk, several printers, and bookshelves for procedures and other operations
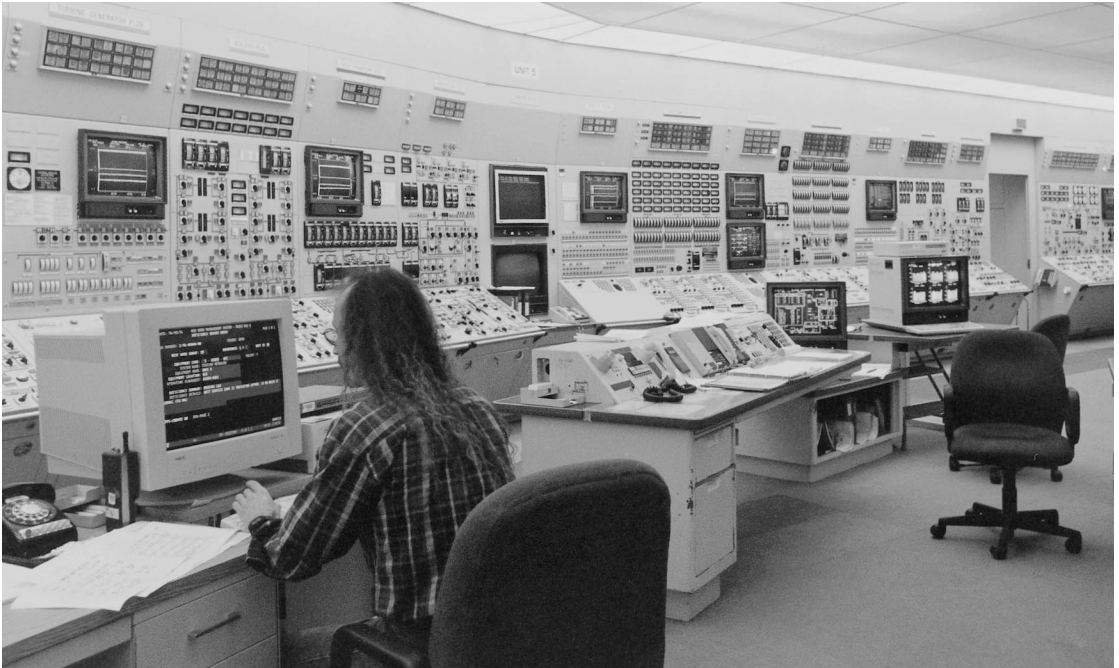
Figure 1. A single unit in the control room.

documents. The control panels are made up of traditional hard-wired meters, strip-chart recorders, and control devices. Alarms (primarily those that are safety related) are presented at the top of the control panels as a series of tiles that light up and provide an audio tone if an alarm condition occurs. There are also several CRTs on each unit. Two CRTs are embedded in the control panels. These are dedicated to presenting textual alarm messages as a chronologically ordered list. In addition, several CRTs are located next to the operator desk for display of plant state information. This plant had a limited number of plant parameters available for display and limited display capabilities (i.e., some physical schematics, trend displays, and bar chart displays). Some instrumentation is located outside the main control room in a separate room behind the unit control panels. These panels are accessed through a doorway next to the unit.

The second plant, Plant B, had a slightly more technologically advanced control room interface. The control room at Plant B is similar in layout to the Plant A control room. Again, four control units are located in a sin-gle, large control room. The operating practices are also similar. The primary differences between the control rooms are in the level of computerization. Plant B has many more plant parameters that are accessible via the plant computer. There are eight CRTs on each unit for display of plant parameter information, of which seven are embedded in the control panels and one is located next to the operator desk. As at Plant A, the alarm system consists of a chronologically organized alarm list presented on the dedicated CRT embedded in a control panel and a series of tiles distributed at the top of the control panels.

Study 1

Our first field study was exploratory in nature. We were primarily interested in answering two questions: What sources of information are available to operators for monitoring? What makes monitoring difficult? In addition, we wanted to get a preliminary understanding of the strategies that operators use to facilitate the job of monitoring.

Method. Two observers independently conducted observations in the control room for

Plant A. Six operators were observed and informally interviewed on an ad hoc basis over the course of 10 12-hr shifts (both day and night). Total observation time was approximately 80 hr. Both observers took notes in the control room but did not communicate with each other during the period of observation. Afterward the observers wrote independent summaries of their findings. At that point the summaries were compared and differences were resolved. A follow-up discussion with some of the operators observed was also used to remove confusions.

Results. As we expected, the control room indicators and alarms were the primary sources of information for monitoring. However, this study readily revealed that operators rely on a much broader and more diverse set of information sources that are used to build and update a mental model of the current plant state. Information sources include shift turnovers, historical information documented in logs and longer-term status binders, and information about certain systems and equipment that is obtained via communication with operators who are located out in the plant (referred to as field operators), written records (check forms), and plant walk-throughs that are periodically conducted by the operators themselves. Each of these is described in more detail in a later section.

In addition to illustrating the diversity of sources of information available for monitoring, Study 1 also gave us some preliminary insights into the factors that make monitoring a plant during normal operations a difficult job. Factors included system complexity and reliability, alarm system design, display and control design, and automation design. Study 1 highlighted that the context for correctly interpreting instruments is complex and constantly changing. As a result, there is no simple answer to the question, How can one tell if the unit is in the desired state? These issues – operator strategies, information sources, and system complexity – are also described in greater detail later in this paper.

This study also documented a number of strategies that operators adopt in order to facilitate monitoring. However, because we observed only six operators, we were not sure to what extent these strategies were representative or idiosyncratic. This question motivated our next two field studies.

Study 2

Method. We circulated a report of our findings from Study 1 (Vicente & Burns, 1995) to 17 operators who had not already been observed. Operators were instructed to focus on the following points in particular: (a) Does each operator perform the strategies documented in the report? (b) Are there strategies that operators use that were not described in the report? (c) Is there anything that was written in the report that is factually incorrect?

Results. Seventeen operators (of the approximately 25 at the plant) reviewed our report. Their comments on the report were passed back to us either in writing (5 operators) or orally (12 operators). Their feedback indicated that the Study 1 report was an accurate portrayal of operator monitoring activities. In addition, operators found our analysis of sources of task complexity and identification of monitoring strategies to be valuable. Comments included minor corrections or clarifications of points made in the report, additional examples of the classes of sources of monitoring complexities or operator strategies defined in the report, and suggestions for changes in design or operations that would facilitate monitoring. Some additional sources of information that can be used for monitoring were also brought to our attention.

Study 3

Method. This study was also conducted in the Plant A control room by the same two observers. They worked independently, observing different operators at different times. Seven operators were observed over eight shifts on various units for a total of approximately 65 hr of observation. None of these operators had participated in Study 1, but some had participated in Study 2. The general method adopted for this round of data collection was the same as that used in Study 1, with the exception that the observers benefited from previous experience. As a result, the observations, though still opportunistic, were more directed than were those in Study 1.

The study's objective was to confirm the classes identified in Study 1, identify additional examples of these classes, and look for new classes that had not been observed in Study 1. One observer used the following items as a reminder of issues to probe: (a) ask "What is your 'regular' monitoring pattern?"; (b) ask "What are the different reasons you might have for monitoring?"; (c) watch for new classes not captured in Study 1; (d) get more examples of the classes identified in Study 1; and (e) get retrospective reports of critical incidents tied to monitoring.

The second observer used a more comprehensive checklist of issues to guide observations, which consisted of (a) plant or crew policy and procedures for guiding monitoring, (b) unit and control room staffing, (c) overview of the control room interface and information resources for the operator, (d) overview of the alarm system, (e) problems with control room technology or operation, (f) use of automation in the interface, (g) how monitoring can become difficult, (h) the motivations for monitoring, (i) operator interactions with the interface to support monitoring, and (j) monitoring strategies. The first six items (a–f) were used largely to confirm previous observations. The primary focus for Study 3 was on the last four items (g–j).

Results. The results of Study 3 confirmed that the findings obtained in the first two field studies could be generalized across operators in Plant A. We collected additional examples of categories that had been identified in the previous report. Also, we described some new categories relevant to knowledge-driven monitoring that had not been identified before.

Collectively, the results from our first three field studies provided a fairly complete description of how operators monitor the plant during normal operations. However, because all three studies were conducted at Plant A, we did not know the extent to which our results generalized across plants. This issue served as the focus of our fourth field study.

Study 4

Method. Study 4 was conducted in the control room of Plant B, which had slightly more computer-based displays than Plant A. Two observers, applying an approach similar to that described in Study 3, each spent approximately 32 hr observing and interviewing operators who were on shift. The two observers were in the control room at the same time but worked independently, each at a different unit. A total of 14 operators were observed and/or interviewed over 10 shifts. Observations and interviews were directed by the lists of issues identified for Study 3.

Results. The observations at the second plant confirmed, and provided additional examples of, the classes of monitoring activity and facilitating strategies identified in the earlier studies. Some of the classes of monitoring activities identified at the first plant were less prominent at the second (e.g., the operators at Plant B reported that they rarely went into the plant to directly observe equipment operations, whereas some operators at Plant A reported that they did so regularly). In addition, some new categories of facilitating strategies were identified (see Mumaw, Roth, Vicente, & Burns, 1996, for more details).

The results of these four field studies provided a set of core findings regarding what makes monitoring difficult and the strategies that operators have developed to facilitate monitoring. Further, we determined that our findings could be generalized across plants with different levels of interface technology. Although differences in technology affected the details of the strategies displayed by the operators, the same classes of facilitating strategies were observed at both plants, establishing the generality of the findings.

The following sections describe these findings in more detail. Note that in presenting the results, we have tried to abstract upward from the specific details of the domain and technologies in order to describe the results at a level that will have meaning beyond the nuclear power plant context.

HOW OPERATORS MONITOR DURING NORMAL OPERATIONS

Contrary to the simple hypotheses that monitoring can be characterized by vigilance, attention, or visual perception, we found that monitoring during normal operations was a

complex, cognitively demanding task that was better characterized as active problem solving than as passive vigilance. The challenge that operators faced was not how to pick up subtle abnormal indications against a quiescent background; rather, it was how to identify and pursue relevant findings against a cognitively noisy background (Woods, 1994). We begin by documenting the sources of information available to operators for monitoring and then explore what makes monitoring difficult. Next, we present the effective strategies that operators developed to maximally exploit the information available and facilitate monitoring.

Sources of Information for Monitoring

We expected that the control room parameters and alarms would be the primary sources of information for monitoring. However, the studies revealed that operators rely on a much broader and more diverse set of information sources.

Shift turnover. An operator arrives in the control room approximately 15 to 30 min before his or her 12-hr shift begins and conducts a shift turnover with the operator being relieved. A common initial activity is a discussion on the state of key parameters, any unusual alarms, jobs completed and jobs outstanding, plans that are active, parameters that need to be monitored more closely than normal, which field operators or technicians are working on which components, what the field operators are aware of, any significant operating memos, and a review of the log (described in the next paragraph). After these discussions, the arriving operator looks at the call-up sheets to see what tests are scheduled for that shift and at the daily work plan that documents upcoming maintenance and call-ups. The operator also reviews the computer summaries and alarms and, at this point, tries to account for every alarm until satisfied that he or she understands why these alarms are in.

Operators are also required to execute a formal panel check procedure to determine whether certain important parameters are in an acceptable state. Some operators were observed to conduct an informal panel walkthrough (systematic review of indicators) before beginning the turnover. Finally, opera-

tors who were not intimately familiar with a unit – for example, operators who had not worked on that particular unit for several weeks – also reviewed the long-term status binder that documents the "quirks" of that unit (e.g., temporary performance anomalies).

Log. The log is a handwritten, chronological record of significant activities (not necessarily abnormal) that have occurred during a shift (e.g., tests completed). This is a short-term record of the history of a unit, as opposed to the longer-term events logged in the long-term status binder. The log is reviewed during the shift turnover, but it can also be consulted during a shift to remind the operator of what had been done on the previous shift or in the last several weeks. The log thereby provides a means by which operators can be aware of the recent status of a unit (e.g., which components are not working, which meters are not working, what is currently being repaired). This provides a valuable context for monitoring and interpreting information on a shift.

Testing. Usually, a number of routine equipment tests are scheduled on every shift. The purpose of these tests is to ensure that backup systems and safety systems are in an acceptable state, should they be required. These tests provide operators with a reason to monitor the status of these systems (e.g., which safety systems are working properly, how quickly they are responding, which meters are working).

Alarm screens. The CRT screens used to display alarms are a salient and frequently used source of information for monitoring. Because the onset and offset of an alarm is accompanied by auditory signals, the alarm screen frequently captures the operator's attention. The operator's response to alarms can vary widely. Some alarms are explained by deficiency reports. Others are caused by maintenance activities. Some are ignored because they are not significant (i.e., they do not affect immediate operating goals) or because there is nothing the operator can do about them (i.e., they cannot be cleared until the unit is shut down). Finally, some alarms motivate a search for more information, typically by sending field operators or maintenance personnel (e.g., control technicians) to conduct additional observations or to perform tests in the plant. It was clear from

our observations that alarm screens play an important role in monitoring plant status.

Control room panels/CRT displays. Obviously, the control room panels (including the alarm tiles) are an important source of information, as are the computer displays that are available for monitoring. Several displays were found to be monitored on a regular basis by virtually all the operators observed.

Field operators. Some parameters and components cannot be monitored from the control room, and control room operators rely on field operators to monitor these parameters and components. Check forms are used periodically by field operators to document the status of certain parameters. These forms are subsequently reviewed by the operator, thereby enabling him or her to indirectly monitor parameters out in the plant. Field operators can also detect problems out in the plant and then bring them to the attention of an operator. Finally, after an alarm, operators sometimes rely on field operators to help diagnose the problem through observations.

Field tour. Periodically, operators take some time to walk through their unit out in the plant. This enables them to maintain a "process feel" by directly observing plant components (e.g., turbine, hydrogen panel, oil purifier panel, boiler feed pumps). Interestingly, this source of information seemed to be relied on more heavily at Plant A (the plant with the older technology) than at Plant B. The difference may be attributed to the more extensive instrumentation of plant state parameters available in the control room in Plant B.

Other units. Other units and other redundant channels on the same unit provide a relatively easy and reliable way of obtaining a referent against which to compare the parameter operators wish to evaluate.

What Makes Monitoring Difficult?

In this section we provide a detailed picture of the context for monitoring, which is influenced by system complexity and reliability, alarm system design, displays and controls design, and the design of system automation.

System complexity and reliability. Each unit consists of thousands of components and instruments. Even though the reliability of each individual component or sensor may be high, equipment failures are bound to occur on a regular basis when there are so many components. Further, some failures can be effectively repaired only when a unit is shut down. Failures of this type (i.e., not essential to the safe and efficient operation of the unit) may persist for a long time before being repaired. For these reasons there are *always* components, instruments, or subsystems that are missing, broken, working imperfectly, or being worked on. Despite this, the unit can still function safely because of redundancy.

Nevertheless, small failures or imperfections have important implications for monitoring. More specifically, they change the way in which information should be interpreted. That is, whether an indicator or set of indicators is normal or abnormal depends heavily on which components are broken, being repaired, or working imperfectly. The same set of indicators can be perfectly acceptable in one context and threaten safety in another. Thus the operational status of the unit's components provides a background, or context, for monitoring. Consequently, effective monitoring depends heavily on an accurate and comprehensive understanding of the current status of plant components and instrumentation. This understanding (mental model) can then be used to derive expectations about what is normal/abnormal, given the current state of the unit. These expectations serve as referents for monitoring.

Two additional aspects of the environment complicate an operator's ability to establish what is normal. First, because there are so many interactions among components, subsystems, and instrumentation, it is difficult to derive the full implications of the current failures to determine what state any particular parameter should be in. Second, this context changes frequently. Some components get repaired, some do not; which components are being maintained can change on a daily basis; new failures can appear, and old problems sometimes go away. It is important that operators be able to track all these changes effectively so they have an understanding of the current context and are able to derive accurate expectations about the alarms and parameter values displayed on the control room panels.

Alarm system design. Operators rely extensively on the alarm system, especially the chronologically organized alarm list presented on the CRTs dedicated to alarm messages, and they work to keep the number of alarm messages low. However, a number of weaknesses of the alarm system compromise the ability of operators to use it effectively.

Many of the weaknesses in the alarm system design arise from the fact that alarm set points are not context sensitive (generally speaking, although there are a few cases for which context is used to suppress alarms). As a result, nuisance (i.e., nonmeaningful) alarms of various types abound. For example, some alarms are always active because the plant is not currently operated the way it was originally intended to be (e.g., because of equipment upgrades). Others appear because a certain component is being repaired or maintained or is not working perfectly.

Nuisance alarms also appear because of a lack of filtering. For example, multiple alarms can appear for the same event and thereby make interpretation more difficult (low-priority alarms are automatically blocked only during an upset). Also, equipment status messages are presented on the alarm monitors even though they do not have the same safety implication.

Perhaps the most frustrating source of nuisance alarms arises from the interaction between sensor variability and rigid alarm limits. Alarms emit one auditory signal at onset and another at offset. If a particular parameter is rapidly cycling above and below the alarm set point, an almost continuous stream of auditory signals is generated.

For all these reasons, the vast majority of alarm messages on the alarm CRTs do not require operator action. Estimates and actual counts at both plants revealed a high percentage of alarms (often higher than 50%) that were not meaningful. Clearly the prevalence of nuisance alarms greatly reduces the informativeness of the alarm system and puts a great burden on operators to distinguish the infrequent alarms requiring actions from those that do not.

Displays and controls design. In both plants we also observed weaknesses in the design of panels, indicators, displays, and controls. The following are general categories of these issues.

(a) Unreliable indicators. There are numerous problems with elements of the interface that too often provide unreliable indicators, which introduces uncertainty for operators. At the lowest level are the simple mechanical elements and lights that make up the control panels. First, at Plant A, indicators about status (e.g., flow paths, pump state) are provided by electromechanical indicators (EMIs). We found that these devices sometimes got stuck in the wrong position, and operators were likely to not trust these indicators. Similarly, at Plant B, LEDs are used to indicate that a variable is in an alarm state. Because these LEDs are difficult to replace, operators usually do not change the burned-out LEDs and lose the information provided by the LED. Also, at Plant A, operators cannot test the light bulbs on the control panel (unlike those in the alarm windows) to see which ones are burned out. Because there are so many bulbs, it is not uncommon for some to burn out. This can create misleading feedback, making it difficult for operators to determine whether an observed anomaly is being caused by a burned-out light or by an actual problem in the unit.

(b) Failed meters. The older, hard-wired meters used at both plants are poor in showing a failed state. One class of analog meters is motor driven, and when these meters fail, the needle remains in the same position – that is, these meters fail "as is," making detection of the failure difficult. Other instruments are designed so that their failed (irrational) value is the same as the low value on the scale. As a result, it is difficult to distinguish a failed sensor signal from a veridical one.

(c) Few emergent features. Because the control panels have so many different types of objects (meters, controls, EMIs, etc.), it would be valuable to have some emergent feature that is easily perceived indicating that an abnormal indicator is present and needs to be located. In general, the control panels do not provide an emergent feature for scanning (e.g., the needles or control handles do not all line up when they are all in the normal state). As a result, the status of these instruments and controls cannot be easily and quickly monitored at a glance or from a distance. Instead, operators must monitor them serially and effortfully, having to recall what the normal position or value of each is and then to determine whether the control or instrument is in that state.

(d) Clear referent values are not always available. Some computer displays (e.g., some bar chart

displays) do not show upper or lower referents for determining whether the current values of the displayed parameters are normal or abnormal. Consequently, these displays require experience, knowledge, and memory to interpret. (Even though these displays may indicate when a value has exceeded the set point, the set point is not available to the operator, so he or she cannot see if the set point is being approached.) We found this also to be true for meters located in areas of the plant outside the control room. Few meters did not require the operator to recall from memory (or a separate document) the referent values.

In some cases, meters are so closely packed together that the indicator and its proximity to the set point are difficult to read from a distance. This is true of the analog meters in both plants that show the neutron overpower values used by the safety systems. Thus a special-purpose monitoring device was retrofitted into the control room to show margin-to-trip with bar charts. This CRT-based display uses color coding and bar length to show when a sensor value is close to the trip point, which makes it easier for operators to extract that information. However, we observed an interesting problem with that display: One of the sensors failed and the corresponding display bar disappeared from the set of 16 bars. In this case, the bar that was lost was the last bar of the set, so it was difficult to see that the bar was not there

(there was no gap). The problem was not picked up until there was a panel check with the panel meter.

Design of automation. There are two types of automated systems: analog and digital. Analog automated systems are governed by individual controllers, the status of which is displayed by an analog meter on the control room panels. Digital automated systems are governed by the computer, and their status is displayed on computer displays that can be brought up on a CRT.

The status of each analog controller is represented by a linear, vertical analog meter, illustrated in Figure 2. There is a green band indicating the set point region for the controller and a red bar indicating the current status. If the red bar is in the green area, the goal is being satisfied. Although it may seem a relatively straightforward task to monitor such controllers, for several reasons the monitoring task is more complex than it appears. First, some controllers are backups and therefore are not controlling. For these controllers, the red bar is not in the green area, even if everything is normal. Therefore, operators must know which controllers are supposed to be in bounds and which are not. To do this, operators need to be aware of what is being done to the plant, what is down, and what should be active,
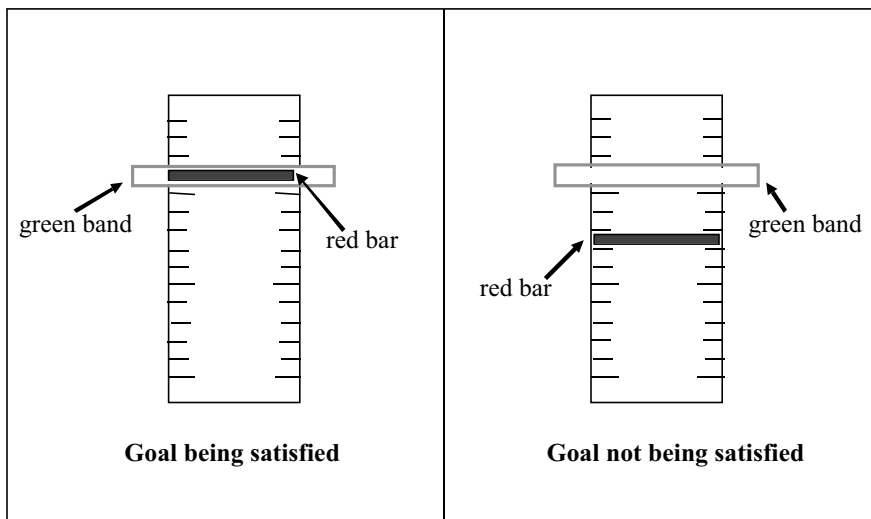

Figure 2. Analog meter showing the status of an analog automatic controller.

given the current state. As mentioned earlier, this is a lot of information to keep in mind.

Second, it is important to distinguish between the actions of an automated system (e.g., trimming valves) and the effects that those actions have on plant parameters (e.g., increased level). From our observations, there appears to be no direct way to monitor the former because that information is not displayed. Thus operators seem to focus instead on monitoring the effects of the automation on plant parameters. This is an important distinction, because if a controller is successfully compensating for a fault (e.g., a leak), no visible signal can be observed in the controller meter; the parameter remains in the goal area. Thus it is possible that the effects of a fault could be masked.

In some cases indirect cues can be used to detect a problem (e.g., decrease in storage tank level). However, this is not possible in all situations (e.g., when makeup water comes from the lake).

In some ways, monitoring of the digital automation is similar to that of the analog automation. The focus is again on monitoring the effects of the automation on plant parameters. Also, alarms go off when the error signal exceeds the threshold, but this alone makes it difficult to perform an effective diagnosis. In addition, however, CRT screens are available to summarize the current status of the various control loops. The problem is that there are more control loops than there are CRTs. Consequently, operators can monitor the effects of only the most important loops (e.g., reactor regulating status). This makes it difficult to monitor, and stay in touch with, the status of all the digital control loops.

Strategies Operators Use to Facilitate Monitoring

The preceding description of the complexity of the nuclear power plant control room environment reveals the impediments to effective monitoring. However, we identified a variety of effective strategies that operators have developed to compensate for limitations in the existing human-system interfaces and to reduce the cognitive demands associated with monitoring. These activities fall into several classes: (a) activities that are designed to enhance information extraction by increasing the salience of important indicators and reducing the background "noise," (b) activities designed to create new information, and (c) activities designed to offload some of the cognitive processing onto the interface (e.g., creating external aids and reminders for monitoring).

These results reinforce and expand on findings in other domains that stress the importance of tailoring strategies developed by domain practitioners to compensate for limitations in system interfaces or domain complexity (e.g., Cook & Woods, 1996; Seminara, Gonzalez, & Parsons, 1977).

Strategies that Maximize Information Extraction from Available Data

Operators have developed strategies that can be used to maximize the information they extract from the plant state data available to them.

Reduce noise. Operators displayed a variety of alarm management activities designed to remove noise so that meaningful changes could be more readily observed. The following are examples of these activities.

(a) Clear alarm printer. At shift turnover, operators clear the printer of all alarms generated on the previous shift. Then an operator can be sure that alarms appearing in the printer happened on his or her shift, thereby facilitating the search for, and organization of, information.

(b) Cursor alarms (i.e., delete the alarm message from the screen before the alarm actually clears, but do not disable it) when they are considered to be unimportant. This strategy keeps the alarm screen uncluttered in order to make new alarm messages stand out.

(c) Disable nuisance program alarms in software. The set points of software-based alarms cannot be changed easily by the operator. Therefore, if a parameter is bouncing in and out of tolerance, a continuous stream of alarms is generated. This defeats the purpose of the alarm and causes the operator to ignore the nuisance alarm (and, understandably, to become extremely frustrated). When this continues for an extended period, the operator may disable that alarm and document this on a paper sticky note that is kept on the side of the CRT. If the alarm is still disabled at the end of the shift, the operator either reconnects it or tells the new operator about it.

Enhance signal. This action increases the salience or visibility of an indicator or piece of

information. For example, we observed operators at one of the plants expanding the *y* axis (the parameter value scale) on a trend plot on the CRT to better monitor small changes in a particular parameter.

Document baseline or trend. This action documents a baseline condition (e.g., at the beginning of the shift) or establishes a trend over a period of time to provide a referent for comparison at a later time. The most common example of this facilitating activity is making hard copies of several CRT displays (more than is required by formal procedures) at the start of a shift. This has the advantage of producing a hard-copy referent of the history of the unit's status so that if unexpected events occur later in the shift, the operator can see (rather than remember) if there was a change from the previous state and, if so, in what way. This historical record thereby provides a valuable context for interpreting subsequent information. Other examples include the following:

(a) Dedicate a CRT-based trend plot to a parameter. It is easier to see change in a parameter when it is plotted with a history (e.g., the last 20 min). In this context change is typically seen through the disturbance of a flat line that stretches across the CRT display. Without the history, the operator must monitor the parameter continuously to see transient changes.

(b) Read the log summaries from the past few days to get continuity. This provides operators with the unit's recent history and establishes a context for the relevance of new information.

(c) Write down parameter values during the panel check procedure. The procedure merely requires operators to check off whether the parameter is within the normal bounds. However, by writing down the exact value on their check sheets, operators can determine how close the parameter is to its boundary limit. If it is close, they monitor that parameter more closely. This strategy serves the general purpose of allowing operators to anticipate problems before they occur.

(d) If an alarm comes in, and there is a reason to suspect that the alarmed parameter may increase (or decrease) even further, the operator may write down the value to monitor for change periodically. The main reason for doing this is there will be no subsequent alarm message if the parameter continues to increase. Thus writing down the parameter and its current value off-loads memory and provides a cue for continued monitoring.

Use other reactor units as referents. Operators mentioned the ability to check equipment and tests with the other three reactor units in the control room as a way of cross-checking information. Operators clearly take full advantage of redundancy throughout the plant using redundant indicators, other channels, and other reactor units to extract the maximum amount of information. This provides them with a means of generating a referent that can be used to determine whether or not a given value is normal.

Exploit knowledge of the plant and the current context to guide monitoring. Operators know that certain types of plant changes (e.g., raising power, refueling) are more likely to cause problems, so they proactively monitor certain parameters more closely (e.g., boiler levels, storage tank level) during those times. Again, this allows them to anticipate problems and catch them at a very early stage.

Operators also know what jobs and/or tests have had problems in the past, either by experience or by looking in the long-term status binder. Using this knowledge, they are prepared to do more careful proactive monitoring of specific parameters that can reveal problems before they become serious.

During the execution of testing procedures, operators always try to understand the intent of a test and not merely follow the procedures in a rote fashion. As a result, they proactively monitor certain parameters to confirm that the test is going as planned. This knowledge-driven monitoring strategy serves several purposes. First, it generates information that can be used to detect errors as soon as they occur. Second, it can help compensate for the limitations in the procedures. Finally, to properly interpret instruments, it is important to know how they can fail. This knowledge can take two forms: knowledge of the internal structure of the instruments, and knowledge of where or when they have failed in the past. Operators use both types of knowledge in monitoring instruments and in interpreting anomalies.

Exploit unmediated indications. Operators may also exploit what we have called *unmediated* indications, labeled as such because there is no sensor-indicator link between the event and the operator. For example, when certain

plant components fail in the open state, a low, rumbling noise can be heard in the control room. Experienced operators use this noise to interpret the state of the reactor unit. Second, the motor-driven meters mentioned earlier make a noise when they move. Therefore, when there is an upset, many of them change simultaneously, thereby providing a salient auditory signal to the operator that something severe has happened (in addition to whatever alarms might come on). A third example is that sometimes the flicker of the lights in the control room is a precursor to problems with the power supply. Experienced operators know this and therefore monitor certain parameters more closely if they notice such a flicker.

Operators may also take full advantage of direct observation of components in the plant (for a fuller discussion, see Vicente & Burns, 1996). The power of direct observation is illustrated by the following incident: Field operators had to go out in the plant to find a heavy-water leak of 50 kg/hr (the shutdown limit). Before they went out, they and the operator went to the sink, with a cup of known volume and a watch, and adjusted the tap flow rate until they created a flow rate of 50 kg/h. As the operator observed, "You might think that 50 kg/h sounds big, that it might be a gusher, but it's not. It was just a trickle." By performing this little experiment, operators acquired a perceptual (rather than symbolic) referent that they could then effectively use for monitoring in the plant.

Seek more information. When an anomalous indicator is observed, operators avoid reacting right away. As one experienced operator put it, "Don't jump in with both feet first." Instead, they consult redundant panel information and/or talk to field operators to establish whether the observed anomaly is being caused by a faulty instrument or by an actual change in plant status. This independent confirmation allows operators to avoid reacting to false indicators, which could lead to problems or trips. Sometimes the problem goes away after a few minutes. Because of the control room hardware reliability issue discussed earlier, many anomalous indicators turn out to be caused by instrumentation problems. This strategy is a way of accommodating this fact.

Strategies that Create Information

Operators have also developed strategies for creating information not normally provided by the existing interface.

Create a new indicator or alarm. We found several instances in which operators modified the interface to create indicators or alarms that did not exist before. Examples include the following.

(a) Operators may change the set points of an alarm after an alarm trips. Because an alarm does not trip again after the value exceeds the set point, operators increase the set point after the initial trip to get a second chance. That way, if the parameter continues to increase, the alarm trips again. Detecting this subsequent auditory signal is much easier than continually having to check on the parameter to see if it has increased even more.

(b) Operators may change an alarm set point on a particular parameter to a temporary value so that an auditory signal occurs when it is time to perform an action (e.g., close a tank valve when the tank has drained to a specified level). Otherwise, they would have to remember to check the meter periodically until its value reached the point at which action was required.

(c) Operators may manipulate alarm set points in ways not intended by alarm designers in order to compensate for the lack of direct information needed for a particular problem. At Plant A, we were told that in some cases operators manipulate alarm set points on some parameters to compensate for the lack of alarms on others. They do this by changing the set point on a parameter that is correlated with the one that they actually want to monitor. Thus the set point on the parameter with the alarm is set at a value that indicates that the parameter of interest (the one without an alarm) has reached an undesirable state. This creates an auditory signal. Without this manipulation, no signal would be given, because the first parameter was not instrumented with an alarm. By generating this new information, operators create an early warning of trouble or signal an important event. Interestingly, this strategy was not observed or reported at Plant B, which was better instrumented.

Determine the validity of an indicator. In some cases, there may be questions about whether or not an important indicator is valid (e.g., because it may conflict with some other information) or a need to determine which indicator to trust. A method used to determine

an indicator's validity is to begin to manipulate plant systems or equipment to see if the indicator responds as expected. We found operators also using alarms to determine the validity of a field operator's report. For certain tasks, field operators need to enter restricted areas that are alarmed. Thus when they enter and exit, the control room (CR) operator gets an alarm. Because CR operators know how long it should take to complete a test or maintenance activity accurately, they can monitor field operators. In one case we observed, the CR operator determined that the field operator could not have completed the assigned work in the time he was in the restricted area.

Strategies to Off-Load Cognitive Demands

We observed operators off-loading memory and attentional demands onto the interface or onto other people.

Create an external reminder for monitoring. Operators may leave the door open on a particular strip-chart recorder to make it stand out from others when it is important to monitor that parameter more closely than usual (e.g., open doors on feedtrain tank levels while blowing down boilers). When several parameters need to be monitored, several strip-chart doors are open, but the chart that is the most critical to monitor is pulled out to distinguish it from the others. This very simple action has an enormous information value. First, it provides an external cue for monitoring. If an operator forgets about the need to monitor and then sees an open door, he is prompted that there must be a reason to monitor that particular parameter more closely. Second, this practice also serves as a cue for others. For instance, when the shift supervisor does rounds, he or she looks at strip charts with open doors. Otherwise, the shift supervisor just passes by them.

Create external cues for action or inaction. External cues are also created to remind an operator about unusual configurations requiring modification in action. As one example, operators may put paper sticky notes on the control room panel to flag unusual indicators. Usually, operators would respond right away to correct such indicators. Thus the sticky note serves as a visual reminder to not react as usual

to the observed signals – that is, the unusual indicator is already known.

As a second example, at Plant B there is a set of analog automatic control devices in a row on one of the control panels. These controllers are normally set on automatic mode. If an operator temporarily changes one to manual, he or she slides the controller out of the panel an inch to indicate that it was intentionally placed in the manual mode.

Employ additional operators. In some cases what is needed to support monitoring are more eyes. An operator who must monitor an indicator closely may be unable to dedicate his or her attention to that task. Several operators mentioned that when workload gets high and there are too many monitoring demands, another operator can be dedicated to a small set of indicators or to the alarm screen.

DISCUSSION

In the introduction we raised two sets of questions. One focused on cognitive processes: How do operators monitor a power plant during normal operations? A second set focused on a methodological issue: How can field studies be used to draw generalizations beyond the specific work setting examined?

We presented a set of cognitive field studies that were performed to examine operator monitoring during normal operations. These studies also addressed the methodological issue by providing an illustration of one approach to gain new insights and obtain generality from field studies.

The field observations revealed that the challenge to monitoring a power plant during normal operations is not how to pick up subtle, abnormal indications against a quiescent background but, rather, how to identify and pursue relevant findings against a noisy background. In this sense the demands of monitoring during normal operations are similar to the monitoring demands during disturbance management (Woods, 1994). Our results showed that operators engaged in myriad setup activities in order to place themselves in a position to be able to detect, recognize, and pursue abnormalities. They worked actively to develop an accurate mental model of current plant

state in order to be able to differentiate expected plant behavior from anomalous findings that need to be pursued. They also worked actively to tailor the control room interface to maximize the detection and pursuit of relevant changes. We found the same pattern of results across the four studies that sampled different operators, different points in time, and plants with different control room interface technology.

These findings are consistent with the existing literature on human monitoring, which indicates the importance of mental models and knowledge-driven processes on guiding monitoring (e.g., Moray, 1986; Sarter, 1996). At the same time, the results advance our understanding of monitoring of complex, dynamic systems by revealing the nature and extent of activity in which operators engage to build an accurate mental model of current plant state and to proactively adapt the interface to maximize detection of meaningful changes.

There Is More to Monitoring...

We described a set of strategies devised by operators to gather relevant information, remove or reduce meaningless changes from the interface, make important information more salient, create information different from that provided by the designers, and off-load cognitive demands onto the environment. Figure 3 summarizes the strategies that are used to facilitate monitoring. (A more detailed model is described in Mumaw et al., 1996.) The figure shows that a major element of operator monitoring is an active, problem-solving process – a form of knowledge-driven monitoring that exploits proactive adaptation of the control room interface.

This proactive adaptation has been referred to as "finishing the design" (Rasmussen, 1986), whereby operators manipulate the interface to make it more informative. This activity helps operators compensate for design deficiencies and for the complex demands imposed by monitoring. It shows that providing skilled operators with responsibility can lead to highly creative and functional forms of adaptation.

In the introduction we put forth three perspectives on monitoring: as vigilance, as selective attention, and as visual perception. Although we found elements of each of these perspectives in our observations, each one fails as a central framework for describing operator behavior. The inability of a single perspective to account significantly for monitoring comes as no surprise. Parasuraman, Molloy, Mouloua,
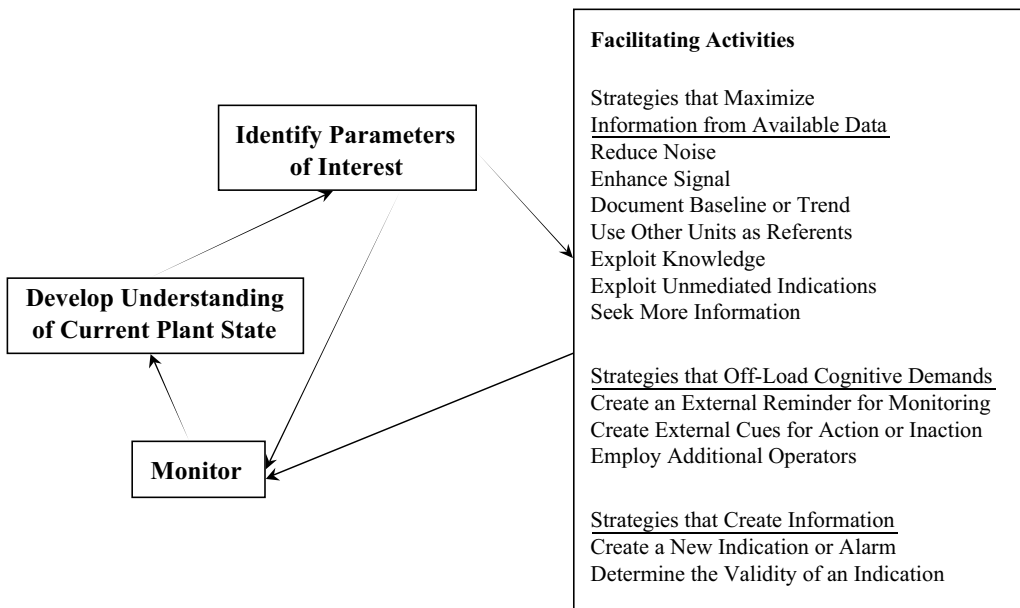

Figure 3. Simple model of knowledge-driven monitoring.

and Hilburn (1996), in their discussion of monitoring automated systems, also identified the complexity of high-technology, high-automation work settings as a determinant of monitoring difficulty. The complex characteristics of the indications and their interface (e.g., unreliability; infrequent, noisy signals), combined with the difficulty in generating knowledge-driven expectations of system behavior, create the need for a complex model of monitoring.

Although vigilance is a relevant issue, our examination of this complex work setting showed that it would be virtually impossible to continuously sample plant indicators in a comprehensive and reliable manner. The number of indicators and displays removes this as a viable option. Furthermore, operators are often occupied with other activities, so monitoring must be time-shared with other tasks. In fact, a nuclear power plant control room must be designed with this in mind (alarms draw operators' attention to parameters that need to be monitored, procedures serves as aids and reminders to check on the status of certain parameters periodically, etc.). Therefore, the lack of continual, comprehensive monitoring is not indicative of negligence or lack of dedication on the part of the operators. Given their task demands, it is the only plausible approach to adopt.

Thus monitoring any complex system involves deciding what is relevant to monitor and what deviation or change is meaningful. The same parameter value can be normal in one context and safety threatening in another. To make these decisions, operators must establish an appropriate context: an up-to-the-minute, comprehensive, and accurate understanding of the unit's operational status, which typically includes equipment status and operational and maintenance activities. Because this context is incredibly complex and constantly changing, maintaining it is not a trivial matter for operators. However, effective operator monitoring relies on this continual mental model development. Our observations suggest that one critical element in acquiring this context is an effective shift turnover (see Edwards & Lees, 1974, and Patterson & Woods, 1997, for a similar finding in other contexts).

On a general level, our observations revealed two different kinds of active, directed monitoring. First, on a regular basis operators monitor a subset of plant parameters because of their importance to plant operations, even in the absence of an alarm. Second, as we have described, it was typical for operators to monitor a subset of parameters more closely because they were tied to the status of an ongoing situation (e.g., maintenance work). Also, on a general level of description, operator monitoring can be handed off to interface technology in stages. The operator can monitor an indicator directly to watch for change, can create a trend plot that reveals change as a deviation from a salient flat line, or can create an alarm that sends a signal when an important deviation occurs. Operators make strategic decisions in assigning limited resources for monitoring.

The field data revealed the central role the alarm screens play in monitoring. Operators seem to rely extensively on the alarm screens to detect problems. In fact, one could even say that to some extent, they use the alarm screens to get an overview of plant status. (It is important to note that this is *not* a function that it was designed to support.)

Finally, we emphasize the contribution of the various informal strategies and competencies that operators have developed to carry out monitoring effectively. Although these strategies are not part of the formal training programs or the official operating procedures, they are extremely important because they facilitate the complex demands of monitoring and compensate for poor interface design decisions. Thus one could effectively argue that the system works well not despite, but *because of,* operators' deviations from formal practices.

Our findings are not attributable merely to poor human factors design of the control rooms. To be sure, there is certainly some room for improvement. However, the second plant we observed is relatively modern, and many of the human factors problems of earlier-generation plants had been resolved. As a result, the difficulties we identified would not have been discovered merely by having a human factors professional go into the control rooms and look at the displays and controls. The important point is that monitoring is difficult because

the context for situation assessment is complex and dynamic. Effective monitoring must therefore be highly context dependent because the very same indicators can be normal or catastrophic, depending on what else is going on in the plant. Moreover, many of these changes in context cannot be anticipated by designers, so operators are required to continually finish the design. This characteristic appears to be a general property of complex sociotechnical systems (Vicente, 1999).

Implications for Interface Design

Our field study findings speak to several issues in interface design. In addition to basic issues described in previous reviews of nuclear power plant interface technology (Seminara et al., 1977), we found little use of emergent features in the design of displays and controls, insufficient filtering of alarms, and a need to make the actions of automation more transparent.

One particular finding that has immediate design implications was the spontaneous manipulation of the alarm system by operators to create auditory cues for conditions for which alarms had not been predefined. This suggests a need to provide operators with a capability to create user-defined auditory cues. Guerlain and Bullemer (1996) described such a system.

Implications for Training

Our field study findings also have important training implications. Operators develop many strategies and acquire a great deal of knowledge on the job that goes well beyond the formal training they receive. Although these strategies are developed and disseminated informally, they could easily be incorporated into a more systematic on-the-job training program. Because of the pressures to keep certification training time from increasing, on-the-job training is preferred.

Another important finding from this field study is the observation that good operators rely extensively on knowledge-driven monitoring instead of rote procedural compliance. This practice allows operators to detect problems before they become significant, to compensate for poor design of procedures, to distinguish instrumentation failures from component failures, and to become better aware (in a deep sense) of the unit's current state. However, current training and licensing programs are based more on procedural compliance than on knowledge-based understanding. The latter set of skills clearly plays an important role in effective monitoring and should therefore play a much larger role in training.

Abstracting Generalizable Results from Field Observation

The series of cognitive field studies we performed illustrates one approach to the problem of how to develop generalizations from field observations that transcend the details of a particular case study. Key features of our approach were

- emphasizing the use of prior research and/or models to guide field observations;

- building a model to organize our results (as embodied in Figure 3);

- assessing the generality of the findings through additional observations – more visits, more plants, plants with different control room interface technology; and

- connecting the results to results of other researchers in other domains of applications.

Our studies used a bootstrap approach (see Hoffman, 1987; Potter, Roth, Woods, & Elm, in press). We began with minimally structured, opportunistic observations and interviews to identify key aspects of the monitoring task and what makes it hard. We then performed a series of additional observational studies using converging methods to test the initial findings and interpretations and to establish generality across individuals and plants with differing technology.

Ultimately the goal is to establish fundamental psychological characterizations that apply broadly across domains. Although our studies examined operator monitoring in nuclear power plants, our findings are likely to generalize more broadly to domains in which operators must monitor multiple complex, dynamic systems. Studies in domains as diverse as space shuttle mission control (Patterson & Woods, 1997), medical operating rooms (Seagull & Sanderson, 1998), steel-

making plants (De Keyser, 1987), and weather-forecasting offices (Hoffman, 1991) provide an indication of similar monitoring demands and similar operator strategies for coping with the domain demands. By establishing commonalities across domains, it becomes possible to build a more fundamental understanding of cognitive performance "in the wild."

ACKNOWLEDGMENTS

REFERENCES

Cook, R. I., & Woods, D. D. (1996). Adapting to new technology in the operating room. *Human Factors, 38,* 593–613.

De Keyser, V. (1987). Structuring of knowledge of operators in continuous processes: Case study of a continuous casting plant start-up. In J. Rasmussen, J. Leplat, & K. Duncan (Eds.), *Human error and new technology* (pp. 247–259). London: Wiley.

De Keyser, V. (1990). Why field studies? In M. Helander (Ed.), *Human factors in design for manufacturability and process planning* (pp. 305–316). Geneva: International Ergonomics Association.

Edwards, E., & Lees, F. P. (1974). *The human operator in process control.* London: Taylor & Francis.

Guerlain, S., & Bullemer, P. (1996). User-initiated notification: A concept for aiding the monitoring activities of process control operators. In *Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting* (pp. 283–287). Santa Monica, CA: Human Factors and Ergonomics Society.

Hoffman, R. R. (1987). The problem of extracting knowledge of experts from the perspective of experimental psychology. *AI Magazine, 8*(2), 53–67.

Hoffman, R. R. (1991). Human factors psychology in the support of forecasting: The design of advanced meteorological workstations. *Weather and Forecasting, 6,* 98–110.

Hutchins, E. (1995). *Cognition in the wild.* Cambridge: MIT Press.

Jordan, B., & Henderson, A. (1995). Interaction analysis: Foundations and practice. *Journal of the Learning Sciences, 4,* 39–103.

Moray, N. (1986). Monitoring behavior and supervisory control. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of human perception and performance.* New York: Wiley.

Mumaw, R. J., Roth, E. M., Vicente, K. J., & Burns, C. M. (1995). *Cognitive contributions to operator monitoring during normal operations* (Tech. Report AECB-2.376.1). Pittsburgh, PA: Westinghouse Science & Technology Center.

Mumaw, R. J., Roth, E. M., Vicente, K. J., & Burns, C. M. (1996). *A model of operator cognition and performance during monitoring in normal operations* (Final Report AECB-2.376.3). Pittsburgh, PA: Westinghouse Science & Technology Center.

Nardi, B. A. (1997). The use of ethnographic methods in design and evaluation. In M. Helander, T. K. Landauer, & P. Prabhu (Eds.), *Handbook of human-computer interaction* (2nd ed., pp. 361–366). Amsterdam: Elsevier.

Parasuraman, R., Molloy, R., Mouloua, M., & Hilburn, B. (1996). Monitoring automated systems. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 91–115). Mahwah, NJ: Erlbaum.

Patterson, E. S., & Woods, D. D. (1997). Shift changes, updates, and the on-call model in space shuttle mission control. In *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting* (pp. 243–247). Santa Monica, CA: Human Factors and Ergonomics Society.

Potter, S. S., Roth, E. M., Woods, D. D., & Elm, W. C. (in press). Bootstrapping multiple converging cognitive task analysis techniques for system design. In J. M. C. Schraagen, S. F. Chipman, & V. L. Shalin (Eds.), *Cognitive Task Analysis.* Mahwah, NJ: Erlbaum.

Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering.* New York: North-Holland.

Roth, E. M., Mumaw, R. J., & Lewis, P. M. (1994). *An empirical investigation of operator performance in cognitively demanding simulated emergencies* (NUREG/CR-6208). Washington, DC: U.S. Nuclear Regulatory Commission.

Sarter, N. (1996). Cockpit automation: From quantity and quality, from individual pilot to multiple agents. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 267–280). Mahwah, NJ: Erlbaum.

Seagull, F. J., & Sanderson, P. M. (1998). Anesthesia alarms in surgical context. In *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting* (1048–1052). Santa Monica, CA: Human Factors and Ergonomics Society.

Seminara, J. L., Gonzalez, W. R., & Parsons, S. O. (1977). *Human factors review of nuclear power plant control room design* (EPRI NP-309). Palo Alto, CA: Electric Power Research Institute.

Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work.* Mahwah, NJ: Erlbaum.

Vicente, K. J., & Burns, C. M. (1995). *A field study of operator cognitive monitoring at Pickering nuclear generating station - B* (CEL 95-04). Toronto: University of Toronto, Cognitive Engineering Laboratory.

Vicente, K. J., & Burns, C. M. (1996). Evidence for direct perception from cognition in the wild. *Ecological Psychology, 8,* 269–280.

Vicente, K. J., Moray, N., Lee, J. D., Rasmussen, J., Jones, B. G., Brock, R., & Djemil, T. (1996). Evaluation of a Rankine cycle display for nuclear power plant monitoring and diagnosis. *Human Factors, 38,* 506–521.

Woods, D. D. (1994). Cognitive demands and activities in dynamic fault management: Abductive reasoning and disturbance management. In N. Stanton (Ed.), *Human factors in alarm design.* London: Taylor & Francis.

Woods, D. D. (1995a). The alarm problem and directed attention in dynamic fault management. *Ergonomics, 38,* 2371–2393.

Woods, D. D. (1995b). Process-tracing methods for the study of cognition outside the experimental psychology laboratory. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsambok (Eds.), *Decision making in action: Models and methods* (pp. 228–251). Norwood, NJ: Ablex.

Woods, D. D., Wise, J. A., & Hanes, L. F. (1982). *Evaluation of safety parameter display concepts* (Final Report NP-2239). Palo Alto, CA: Electric Power Research Institute.

Randall J. Mumaw is Human Factors Specialist in the Aviation System Safety organization at Boeing Commercial Airplane, Seattle, Washington. He received a Ph.D. in cognitive psychology from the University of Pittsburgh in 1985.

Emilie M. Roth is Principal Scientist of Roth Cognitive Engineering, Brookline, Massachusetts. She received a Ph.D. in cognitive psychology from the University of Illinois at Urbana-Champaign in 1980.

Kim J. Vicente is founder and Director of the Cognitive Engineering Laboratory at the University of Toronto, where he is also a professor of mechanical and industrial engineering and of biomaterials and biomedical engineering. He received a Ph.D. in mechanical engineering from the University of Illinois at Urbana-Champaign in 1991.

Catherine M. Burns is an assistant professor of systems design engineering at the University of Waterloo, where she directs the Advanced Interface Design Laboratory. She received a Ph.D. in 1998 in mechanical and industrial engineering from the University of Toronto.