Review of Alarm Systems for Nuclear Power Plants

Kim J. Vicente

CEL 96-04

Cognitive Engineering Laboratory Department of Mechanical & Industrial Engineering University of Toronto
5 King's College Rd. Toronto, Ontario, Canada M5S 3G8
Phone: (416) 978-7399 Fax: (416) 978-3453
Email: benfica@mie.utoronto.ca URL: http://www.ie.utoronto.ca/IE/HF/CEL/homepage.html

Director: Kim J. Vicente, B.A.Sc., M.S., Ph.D.

The Cognitive Engineering Laboratory (CEL) at the University of Toronto (U of T) is located in the Department of Mechanical & Industrial Engineering, and is one of three laboratories that comprise the U of T Human Factors Research Group. CEL began in 1992 and is primarily concerned with conducting basic and applied research on how to introduce information technology into complex work environments, with a particular emphasis on power plant control rooms. Professor Vicente's areas of expertise include advanced interface design principles, the study of expertise, and cognitive work analysis. Thus, the general mission of CEL is to conduct principled investigations of the impact of information technology on human work so as to develop research findings that are both relevant and useful to industries in which such issues arise.

Current CEL Research Topics

CEL has been funded by Atomic Energy Control Board of Canada, AECL Research, Alias|Wavefront, Asea Brown Boveri Corporate Research - Heidelberg, Defense and Civil Institute for Environmental Medicine, Japan Atomic Energy Research Institute, Natural Sciences and Engineering Research Council of Canada, Rotoflex International, and Westinghouse Science & Technology Center. CEL also has collaborations and close contacts with the Mitsubishi Heavy Industries and Toshiba Nuclear Energy Laboratory. Current CEL projects include:

- Studying the interaction between interface design and adaptation in process control systems.
- Understanding control strategy differences between people of various levels of expertise within the context of process control systems.
- Developing a better understanding of the design process so that human factors guidance can be presented in a way that will be effectively used by designers.
- Designing novel computer interfaces to display the status of aircraft engineering systems.
- Developing and evaluating advanced user interfaces (in particular, transparent UI tools) for 3-D modelling, animation and painting systems.

CEL Technical Reports

For more information about CEL, CEL technical reports, or graduate school at the University of Toronto, please contact Dr. Kim J. Vicente at the address printed on the front of this technical report.

INTRODUCTION

This report describes the state-of-the-art of human factors (HF) issues in alarm systems design for nuclear power plant control rooms. It is divided into two parts. The first part describes the important work recently completed in this area by the United States Nuclear Regulatory Commission (USNRC). The second part provides a review of the literature. To anticipate, the findings show that many different alarm system design techniques have been proposed, but that very little empirical research has been conducted. As a result, we simply do not have a very good understanding of the relative benefits and disadvantages of the techniques being proposed in terms of their impact on human performance and plant safety. Despite this, new alarm systems are being designed by nuclear power plant vendors world-wide. Therefore, alarm systems design is an area where there are a number of pressing issues which are in dire need of research.

USNRC POLICIES

Information regarding USNRC policies was obtained by consulting a number of technical reports (cited below), and by a telephone conversation with Dr. John M. O'Hara of Brookhaven National Laboratory (BNL) (June 11, 1996). Dr. O'Hara has conducted a number of projects for the USNRC which are very pertinent to the HF of annunciation systems. As a result, he was an extremely valuable resource.

Background

Before we describe the work recently conducted by BNL, it is important to briefly provide some background context regarding the regulatory environment in the US. Two points in particular are worth making.

First, it is important to point out the distinction between guidance and regulation. Guidance constraints are recommended, whereas regulation constraints are mandatory. For example, NUREG-0700 (USNRC, 1981), which provides a variety of HF guidelines for traditional CRs, is a guidance document. In contrast, the mandatory installation of a Safety Parameter Display System (SPDS) in all nuclear plants in the US after the Three Mile Island accident is an example

of a regulatory requirement. A second example is the requirement that all plants conduct detailed control room design reviews (CRDRs) after Three Mile Island. As far as we know, these are the two primary requirements pertinent to HF issues that have regulatory, as opposed to guidance, status in the USNRC.

Second, it is worthwhile discussing the procedure for making design modifications to CRs that are currently in operation, since this procedure gives some insight into how rigorously such changes are evaluated. The most relevant document for this purpose is US Code of Federal Regulations 50.59, which states that the proposed design change must be evaluated by the utility against its standard safety analysis report (SSAR). If this evaluation reveals no unresolved safety questions (USQs), then the utility is allowed to introduce the proposed change without first obtaining USNRC approval. Because utilities may lack awareness of the potential impact of control design changes on system safety, it is possible for a CR design change that is significant from a HF point of view to be implemented by the utility without first obtaining USNRC approval. This is known as a 50.59 change. Each year, a list of 50.59 changes must be submitted to the USNRC. However, the extent to which these lists are audited is not known. Given the limited resources of the USNRC, it would not be surprising if these changes do not receive the attention they deserve.

Given this background, we can now move on to describing the guidance documents that have been used by the USNRC in the past.

Older Guidance Documents

There are two older guidance documents that have traditionally been used by the USNRC to review and evaluate CR designs, including annunciation systems. These are NUREG-0700 (USNRC, 1981) and Chapter 18 of NUREG-0800 (USNRC, 1984). We will briefly point out the role of these two documents to provide a context for the more recent guidance documents to be described in the next subsection.

The original version of NUREG-0700 (USNRC, 1981) contains some guidance that is pertinent to the design of annunciation systems and to other CR design issues as well. However,

that document only addresses HF issues associated with traditional technology (i.e., analogue, hard-wired instrumentation). Consequently, it is of very limited value in reviewing advanced control rooms (ACRs) or CR upgrades based on computer technology.

Chapter 18 of NUREG-0800 is also a relevant document (USNRC, 1984). It provides a standard review plan for HF issues that the USNRC can use to conduct a review of regulation issues, such as SPDS and CRDRs. This document also provides guidance on how the USNRC should review a SSAR submitted by a utility. This review takes the form of a Safety Evaluation Report (SER) written by the USNRC, evaluating the utility's SSAR. Chapter 18 of NUREG-0800 thereby provides guidance on how the USNRC should deal with the HF issues in a SER. However, as it stands, NUREG-0800 suffers from the same limitations as NUREG-0700. It too is based on traditional CR technology, and is not really relevant to the evaluation of ACRs.

In an effort to remedy this situation, the USNRC contracted BNL to conduct a number of studies to bring this guidance up to date, to the extent possible. These studies resulted in several technical reports, which will be described next.

Newer Guidance Documents

There are 4 documents that have recently been written by BNL for the USNRC that are very pertinent to the HF of annunciation systems:

a) NUREG-0711, "Human Factors Engineering Program Review Model"

b) NUREG/CR-5908, "Advanced Human-System Interface Design Review Guideline"

c) NUREG/CR-6105, "Human Factors Engineering Guidance for the Review of Advanced Alarm Systems"

d) NUREG-0700, Revision 1, "Human-System Interface Design Review Guideline".

Each of these documents will now be briefly discussed in turn.

NUREG-0711 (USNRC, 1994) is perhaps the most important contribution of all, describing a HF program review model (PRM). The PRM is based on the belief that it is necessary to review the design process in addition to the usual process of reviewing the final design product. Consequently, the PRM provides a set of process criteria that can be used to evaluate the process

by which a design is developed, and in particular, the way in which HF issues have been incorporated into the design life cycle. The criteria set out in this document are very broad, being based on a systems approach to HF design. The entire plan consists of 10 elements, each dealing with a key set of HF issues:

1. Program Management - this element of the PRM specifies the authority and responsibility that the HF team is to have within the design organization, if a successful design is to result. In addition, this element specifies in detail the inter-disciplinary expertise that the HF team should include to implement a systems approach to design.

2. Operating Experience Review - this element ensures that designers are aware of the lessons that have been learned in the past, and that those lessons are incorporated into the design being proposed.

3. Functional Requirements Analysis & Function Allocation - this element ensures that designers identify the functions that are necessary to achieve plant objectives, and then allocate those functions according to accepted HF principles. The process advocated is very similar to that used in the US military for many years.

4. Task Analysis - this element of the PRM requires that a task analysis be conducted and documented to inform subsequent design decisions.

5. Staffing - this element requires that designers identify the number and qualifications of the personnel that are required to operate the plant, based on the analyses conducted in the previous PRM elements.

6. Human Reliability Analysis - this element requires designers to analyze the impact of HF issues on plant safety using traditional, quantitative risk analysis techniques.

7. Interface Design - this element ensures that designers have translated the results of the aforementioned analyses into design criteria and implications. In addition, the resulting interface design should be consistent with relevant HF design guidance documents.

8. Procedure Development - this element of the PRM ensures that the design of procedures is based on the same design process that is used to determine other systems components (e.g., interface), relying on a common set of analyses. This will ensure that procedures are well integrated with the remainder of the design.

9. Training Development - this element requires designers to identify the competencies that are required of staff to operate the plant effectively. In addition, an integrated training program that can impart these competencies should be designed according to HF principles and practices.

10. Verification and Validation - the final element of the PRM ensures that the design has been properly evaluated so as to conform to design objectives and accepted HF principles. This involves not only conformance with HF guidelines, but also dynamic evaluations using realistic scenarios and representative testing conditions.

The PRM is intended to be used in the licensing of ACR designs, and to review the SSARs submitted by utilities. The original version was published in July, 1994, but a new, slightly revised version of the document will be available soon.

Another relevant document is NUREG/CR-5908, whose purpose is to compile together the available HF guidance that is pertinent to computer-based interfaces, thereby addressing the

limitations associated with NUREG-0700. NUREG/CR-5908 actually consists of two volumes. Volume 1 (O'Hara, 1994) provides a detailed discussion of the gaps in the available guidance for evaluating ACRs and retrofits based on computer-based technology and proposes a methodology for addressing those gaps. This methodology consisted of a review of existing HF guideline source documents and a procedure for identifying a set of guidelines that are pertinent to the specific needs of nuclear power plant CR design. The resulting guidelines were evaluated through peer-review, field testing, and a workshop attended by experts in the area. Volume 1 ends by providing a prioritized list of areas where the guidelines available are either weak or non-existent.

Volume 2 of NUREG/CR-5908 (O'Hara, Brown, Baker, Welch, Granda, & Vingelis, 1994) describes the guidelines that were compiled using the methodology described in Volume 1. The guidelines were organized according to the following sections:

Section 0. High-level Human-System Interface Design Evaluation Principles

Section 1. Information Display

Section 2. User-System Interaction

Section 3. Process Controls and Input Devices

Section 4. Alarm Systems

Section 5. Analysis and Decision Aids

Section 6. Inter-Personnel Communication

Section 7. Workplace Design

In addition, a set of procedures for using these guidelines to conduct ACR design reviews is also described. These procedures are integrated with the global design process specified by the PRM in NUREG-0711.

Interestingly, one of the outcomes of the work done for NUREG/CR-5908 was that the guidelines in source documents pertinent to alarm systems was weak. As a result, a decision was made not to include any guidelines in that section of the document. To address this substantial gap, a project specifically geared towards the development of guidance to support the USNRC

review of advanced alarm systems was undertaken. The results of this project are documented in NUREG/CR-6105 (O'Hara, Brown, Higgins, & Stubler, 1994). That document describes the methodology that was used to develop the alarm guidelines, presents the guidelines themselves, and provides a procedure for the review of an alarm system.

The fourth and final guidance document referred to at the start of this subsection is an updated version of NUREG-0700 (Rev. 1) (O'Hara, Brown, Stubler, Wachtel, & Persensky, 1996). The objective of this document is to take the place of the original NUREG-0700 by incorporating the latest research findings that are relevant to the design of ACRs based on computer technology as well as guidelines that are most pertinent to traditional CRs. NUREG-0700 (Rev. 1) is best viewed as an integration document that incorporates the results of all of the newer documents that have been described in this subsection. Thus, in addition to including some of the contents of the original NUREG-0700, NUREG-0700 (Rev. 1) also includes the more recent guidance developed in NUREG-0711 (the HF PRM), NUREG/CR-5908 (Advanced Interface Guidelines), and NUREG/CR-6105 (Advanced Alarm Systems Guidelines). The document is ultimately intended to be used for the following purposes: review of a CR for a new plant; review of voluntary upgrades or modifications made to existing CRs; USNRC staff investigation of events involving human performance.

The document consists of two parts. The first part provides a set of criteria that the USNRC can use to evaluate an applicant's (i.e., utility's) own Human-System Interface design review process. The goal is to ensure that the applicant has systematically identified and resolved the human engineering discrepancies that could adversely affect the safety of the plant. This is accomplished by ensuring that the applicant followed, and documented the results of, a systems approach to the HF issues that are pertinent to the design, modification, or issue being evaluated. Part 2 of NUREG-0700 (Rev. 1) contains a set of detailed HF guidelines that can be used to evaluate both advanced and conventional CRs. A number of sections, corresponding to different HF issues, are presented, the most pertinent being section 4 which contains alarm review guidelines (taken from NUREG/CR-6105). This section of the document is 39 pages

long. A number of other sections in the document are also pertinent, especially those which contain general guidelines (taken from NUREG/CR-5908) that are relevant to any controls and displays, including those that are part of an alarm system.

An innovative feature of NUREG-0700 (Rev. 1) is that, in addition to being available in a hard-copy version, the guidelines in part 2 have also been incorporated into an interactive, computer-based review aid that is intended to make it easier to: access and review the guidelines; compile individual guidelines for a specific review; and integrate new guidelines when they become available.

While the work conducted by O'Hara and colleagues at BNL for the USNRC is very impressive in scope as well as depth, it is inherently limited by the state of knowledge of the field. More bluntly, it is difficult to develop guidelines for questions whose answers simply are not yet known. Particularly with the topic of alarm systems, the state of knowledge is such that the research findings available are meagre with respect to the broad range of questions that designers face. Thus, even the new version of NUREG-0700 (Rev. 1) does not provide a great deal of guidance for the design of advanced alarm systems, primarily because many important issues remain to be investigated.

Future Work

In recognition of this fact, the USNRC and BNL have developed a research plan to answer some of these outstanding questions (O'Hara, Wachtel, & Persensky, 1995). Three experiments on alarm systems are to be conducted at the OECD Halden Reactor Project in Norway. Experiment 1 will evaluate the effect of display type on human performance, independently of any alarm processing techniques. Experiment 2 will evaluate the impact of alarm processing techniques on human performance. In particular, the effect of alarm reduction techniques and differential availability of alarm processing results will be investigated. Finally, Experiment 3 is designed to evaluate the impact of alarm generation (i.e., higher-order alarms derived from lower-level alarms) on human performance. All experiments will include scenarios that range in complexity. The current estimate is for these experiments to start in the fall of 1996. Regardless

of the specific outcomes obtained, these experiments will represent an important contribution to the current impoverished level of understanding of the impact of advanced alarm systems on human performance in nuclear power plants.

LITERATURE REVIEW: HUMAN FACTORS OF ALARMS

Scope

This section of the report presents the results of a literature review of HF issues in the design of annunciation systems for nuclear power plants. (Note that the terms "alarm" and "annunciation" will be used interchangeably in this section, reflecting the usage observed in the literature.) Because of limited resources, we were not able to conduct a thorough search of the literature and were not able to consult an extensive number of original source articles. We tried to compensate for this by relying heavily on existing secondary reviews of the literature, complemented by a selective reading of some key primary sources. We found the following secondary sources to be of particular use, and have borrowed extensively from them: O'Hara and Brown (1991); Lupton, Lapointe, and Guo (1992); O'Hara et al. (1995); and especially, O'Hara and Brown (in press). Given the process that we followed, the literature review is not intended to be complete nor highly original. Some of the topics which are explicitly beyond the scope of this review are: operator modification of alarm limits, alarm silencing and acknowledgement controls, alarm response procedures, physical layout, hardware and software reliability, generic lower-level display and control issues (see O'Hara & Brown for a treatment of some of these issues). Despite these limitations, the literature review is intended to present a faithful picture of the most significant techniques and findings pertinent to the design of advanced annunciation systems for nuclear power plants.

The review is divided into five sections, which are a subset of the categories adopted by O'Hara and Brown (in press). The first section deals with the issue of purpose, investigating the different roles that an annunciation system can play in CR operation. The second section deals with the criteria by which alarm states are defined. The third section reviews various advanced techniques that have been proposed for processing alarms. The fourth section deals with the

issue of how alarm information is presented to operators. The fifth section lists, and briefly describes the trends exhibited by, the new advanced alarm systems that are being designed and developed by various vendors in the nuclear industry worldwide.

Purpose: What is the Role of an Annunciation System?

The maturity of any area of research can be evaluated by determining if there is a consensus on defining basic terms. In this case, this translates into whether there is an agreement in the literature as to the purpose of an alarm system. We have found that it is useful to distinguish between the purposes that an annunciation system is intended to fulfill by designers, and the purposes that it actually fulfills for operators, in practice. The former will be referred to as designed purposes whereas the former will be referred to as the operational purposes. Each of these will be reviewed in turn.

Designed purpose. There is a surprising lack of consensus among designers as to what role an annunciation system should be designed to serve in a CR. A few examples will serve to illustrate this point. For example, O'Hara and Brown (1991) state that an alarm system should:

- detect a specific off-normal condition
- provide annunciation of plant status
- indicate specific entry conditions to the emergency operating procedures (EOPs)
- indicate which branch points should be followed in an EOP.

Lupton, Lapointe, and Guo (1992) provide a different definition of the role of an annunciation system, namely to alert operators to changes in plant status. Under this broad umbrella, they include the slightly more specific roles of providing information to monitor the plant, interpret plant events, and take compensatory actions. This is a very broad definition of an alarm system since it includes providing support for virtually all aspects of the operator's job (excluding perhaps administrative tasks).

Easter and Lot (1992) provide a third view of the objectives of an alarm system:

- to alert operators to abnormalities
- to convey information about future consequences of an abnormality.

Compared to the first two, this definition puts greater emphasis on the potential predictive informativeness of an annunciation system.

Woods (1995) sees the purpose of an alarm system as being to direct the attention of operators. As such, it can serve to interrupt, prompt, or remind operators. This definition differs from the others reviewed above in that it describes the role of an alarm system from a psychological perspective rather than from a task demands perspective.

Baker, Hollnagel, Marshall, and Øwre (1985) identified 4 purposes for an alarm system: to alert, inform, guide, and confirm. This definition is similar to that of Lupton et al. (1992) in that it seems to encompass the entire range of operator decision tasks (see Rasmussen, 1976).

Reiersen, Marshall, and Baker (1987) provide a sixth definition of the purposes of an alarm system:

- to alert operators to a disturbance
- to guide operators to the disturbed area
- to help operators locate the disturbed variables.

This definition differs from some of the others since it is focused exclusively on off-normal situations.

Yet another viewpoint is provided by Stanton (1994b) who lists a range of roles that an alarm system can serve:

- to display an unexpected change in system state
- to display changes in system state
- to attract the operators' attention
- to arouse the operators
- to change the operators' mental state

This list of purposes is interesting because it makes some distinctions that are not explicitly captured in the other definitions (e.g., expected vs. unexpected changes in system state).

Finally, some researchers have even suggested that the alarm system should also be responsible for diagnosing the state of the plant, and in some cases even response planning,

thereby integrating fault detection with fault diagnosis and compensation into a single automated system. This viewpoint was prevalent in the Disturbance Analysis System (DAS) work conducted in the 1980s (see Lees, 1983; Bray, 1989; Kim, 1994 for reviews).

This set of definitions clearly shows that there is a lack of consensus in the literature as to what an alarm system is for (cf. Stanton, 1994b). While there is certainly some common ground across the various definitions, there are also significant differences as well. Some of the important areas of disagreement include:

- does the alarm system include normal status, as well as off-normal occurrences?
- does it include expected as well as unexpected indications?
- does it, by itself, explicitly support decision making and response planning activities?

Even worse, the same authors provide different definitions in different papers (compare Baker et al., 1985 with Reiersen et al., 1987)!

Usher (1994) brings this significant problem to the fore when he asks whether an alarm is:

- a part of a control desk (i.e., the annunciator tiles themselves)?
- a system state (i.e., an off-normal event)?
- an audible noise (i.e., the sound emitted by the CR interface)?

The problem seems to be that the label "alarm system" is used in such a way that it confounds a number of different characteristics that are, in principle at least, conceptually independent. It is important to untangle these characteristics so that the full range of design possibilities is clearly revealed. Before we attempt such a critical analysis, however, it is worthwhile examining the operational purposes for which annunciation systems have been used.

Operational purposes. There are relatively few well-documented field studies investigating how NPP operators actually use alarm systems. A recent exception is the work by Vicente and Burns (1995, 1996) which documents the strategies that operators at Pickering B use to monitor the state of the plant. Although this study was not focused exclusively on alarms, the results show that the alarm system plays a very prominent role in operator cognitive monitoring. Interestingly, many of the results obtained by Vicente and Burns (1995, 1996) are consistent with

a field study conducted years earlier by Kragt and Bonten (1983) in a fertilizer plant, which was specifically focused on operators' use of a conventional alarm system.

Both of these studies reveal that the alarm system plays a very important, multi-faceted role in helping operators monitor plant status. Instead of continuously monitoring the plant via a large number of instruments, operators frequently rely on the alarm system to bring their attention to goal-relevant events in ways that were not anticipated by designers. Moreover, the alarm system is used for a myriad of purposes that have nothing to do with alarms, in the sense of an abnormal event. Kragt and Bonten (1983) summarize this by stating that the alarm system was used primarily "as a monitoring tool and not as an alarm system requiring action" (p. 586). Because many of these uses were not intended and not systematically supported, this may compromise the alarm function of the system.

Collectively, the results of Vicente and Burns (1995, 1996) and Kragt and Bonten (1983) reveal that alarm systems are used by operators for the following purposes:

- to detect disturbances that occurred in the process, and if so, what kind (the traditional designed role of an alarm system)
- to determine if certain components or subsystems started up automatically (e.g., backup system, or a safety system), when the process required it
- to determine if disturbances disappeared of their own accord (usually by waiting to see if the alarm would go away by itself)
- to determine if the operators' actions had the desired effects
- to confirm that an event (disturbance or action) that the operator expected to happen did indeed occur
- to evaluate whether a field operator had performed his task properly (and thus, whether the field operator was trustworthy)
- to determine if a field operator had acted on his own initiative without informing the CR operator
- to determine the location of a field operator

- to determine when it is time to perform a certain action
- as an external reminder to monitor a certain parameter closely
- to determine if a parameter is going even further out of its normal operating range, after it has already triggered an alarm
- to create alarms for parameters that do not have designed alarms associated with them.

Examples of each of these uses can be found in Kragt and Bonten (1983) and Vicente and Burns (1995, 1996). These operational purposes are very important because they differ considerably from most, if not all, of the definitions of the designed purposes of alarm systems described in the previous subsection.

Several preliminary conclusions can be drawn as a result. First, alarm systems are used for purposes that were not anticipated by designers and that the alarm system was not designed to support. Second, alarm systems are used for many purposes that are not associated with off-normal events. In fact, one of their primary operational purposes is to monitor the plant during normal operation, to help operators update their situation awareness of the plant (Mumaw, Roth, Vicente, & Burns, 1996). (It should be noted, that both of these studies observed the usage of alarms over short periods, and thus the results obtained are more pertinent to the day-to-day usage of an alarm system rather than the usage during serious plant failures.)

Critical analysis. The preceding subsections have argued that there is a lack of consensus in the literature as to what the designed purposes of an alarm system should be, and that the operational purposes for which alarm systems are actually used in practice go well beyond their designed purposes. How can this set of conflicting findings be explained? We believe that the answer to this question can be found by unraveling a set of characteristics that have sometimes been confounded in the alarm literature. This can be accomplished by drawing the following set of conceptual distinctions.

One important distinction is that the role of an alarm system can differ as a function of the technology upon which it is based (O'Hara & Brown, 1991). More specifically, traditional alarm systems tend to be independent, whereas alarms that are being designed and proposed for ACRs

tend to be more integrated with other parts of the CR interface (e.g., overview display, individual process status displays). With this trend towards integration, there is a commensurate trend towards increasing the number of purposes that an alarm system can fulfill. This could explain why different authors posit different definitions for the role of an alarm system. This explanation does not completely solve the problem, however. Because of the trend towards integration, the distinction between an alarm system and the rest of the CR interface becomes fuzzier and fuzzier. This makes it very difficult, if not impossible, to talk about an alarm system in isolation.

A second important distinction is that the role of an alarm system is different in normal operations (including minor faults) than it is in accident situations (Marshall & Baker, 1994). As some of the definitions of designed purposes above illustrate, the focus has increasingly been on trying to design alarms to support operators under accident situations. This trend was largely prompted by the lessons learned from the Three Mile Island incident, where the well-known problems associated with an alarm "avalanche" reached an inauspicious peak. As a result, many techniques have been devised to help operators interpret off-normal events (mainly by prioritizing and reducing the number of alarms, see below). However, under normal operations, alarms are used for a very different set of purposes. This is clearly illustrated by the list of operational purposes for which alarms have been used in practice, described above. The main potential problem under normal operations is not alarm flooding, but rather the lack of salient, informative (i.e., diagnostic) indications. Evidence for this assertion comes from the creative strategies that operators adopt to extract information from existing alarm systems (Vicente & Burns, 1995). Thus, it seems highly likely that the type of alarm design techniques and features that are required to support normal operations and accidents are quite distinct.

The specific use of annunciation systems to diagnose the status of the plant in the case of an abnormal event deserves a special mention. Automated DASs were proposed by some, after Three Mile Island, as one of the purposes that an alarm system should fulfill. However, experience has shown that such systems have severe limitations associated with them (see Bray, 1989 and Kim, 1994). Some of the reasons for this failure can be found in the operational

realities of NPPs (Vicente & Burns, 1995). Each day, one or more plant components are being tested, are being maintained, are working imperfectly, or are not working at all. As a result, a given set of symptoms can be perfectly normal or absolutely catastrophic, depending on the current state of the plant. Automated systems have a very difficult time dealing with this context-dependency. Because the local context can literally change on a daily basis, designers cannot possibly anticipate, and therefore design for, these contingencies. Instead, they must be confronted, at the time, by the operators since they are the only ones who have access to the unique details specific to the local context. As a result, it is not surprising that the DAS efforts that attempted to automate diagnosis and response planning functions encountered significant implementation difficulties. It is more prudent to develop less ambitious support systems that will provide human operators with the information they need to accurately and efficiently diagnose the state of the plant in the event of an abnormality, rather than allocating such a task to a computer (cf. Lees, 1983, p. 671).

A third crucial distinction is between the auditory and visual modalities. Because of the way in which annunciation systems have been designed in the past, there is a very strong tendency to associate the auditory modality with alarms. But clearly, the auditory modality can be used to communicate information that has nothing to do with an abnormal situation. This design possibility does not seem to have been appreciated in the past.

An excellent example illustrating the potential value of this idea is the design concept recently proposed by Guerlain and Bullemer (in press). Based in part on the results of Vicente and Burns (1995), Guerlain and Bullemer developed a concept for helping operators monitor plant status. This concept, known as user-initiated notification, is similar to existing alarm systems in that it relies on the auditory modality. As a result, it has the advantage that the interface can grab the operator's attention, rather than requiring the operator to periodically monitor the interface. The novel feature of user-initiated notification, however, is that the operator specifies the event that will trigger an auditory signal. For example, if an operator wants to know when a certain vessel has been filled, so that he can turn off a pump, he can tell the

interface to provide him with a signal when the level of the vessel has reached the desired value. Then, rather than having to periodically monitor that parameter (and potentially forget to do so), the operator can merely wait for the system to indicate to him that the event he specified has occurred, thereby liberating him to perform other duties in the meantime. User-initiated notification may thereby give operators more flexibility and control in specifying the information needs that are relevant to the local (unpredictable) context, thereby addressing the types of limitations encountered with DASs (see above). At the same time, this design concept also takes advantage of the benefits of the auditory modality. In fact, such an interface could be used for many of the operational purposes for which alarm systems are currently being used (Guerlain & Bullemer, in press). The important difference, however, would be that the interface would be explicitly designed to support these purposes. Currently, these purposes have been informally adopted by operators and therefore are subject to error (e.g., forgetting to return an alarm setpoint back to its nominal value after changing its value to serve as a reminder for action).

Summary. The results of this critical analysis can be summarized as follows:

1. The purposes of an alarm system in a retrofit of a traditional CR will differ from those of an ACR. It is important to specify the type of context a particular design is intended for, and to design the purposes accordingly. The entire CR interface should be viewed as an integrated system for normal operation management, fault management, and outage management. The annunciation system is only one of the constituent subsystems of the overall system. Taking this approach allows one to recognize that other (possibly novel) subsystems can better serve some functions previously served by traditional annunciation systems.

2. Different interfaces will be required to support operators under normal operations than under abnormal operations. Systems for the former mode are monitoring tools, not

alarm systems (in the sense of detecting plant accidents). The characteristics of these interfaces will need to be different, although both can rely on the auditory modality.

3. Experience has shown that fault diagnosis should not be automated as part of an advanced alarm system. Instead, it is more prudent to limit the role of an alarm system to that of an information provider to a human operator who is responsible for making decisions with respect to fault diagnosis.

Alarm State Definition

A very basic question which apparently has not received a great deal of attention in the literature is what criteria should be used to define an alarm state. Traditional alarm systems have been based primarily on the single-sensor-single alarm philosophy (i.e., each alarm state is defined as an upper and/or lower limit on a particular sensor), but this philosophy does not specify which parameters should have alarms associated with them. In traditional alarm systems, alarms are a very heterogeneous set of plant states, including: passage doors being open; actuation of automatic safety systems; individual parameters going out of their nominal range; and, large-scale accidents. Despite the very large number of alarm states, there are still many parameters that do not have alarms. In some cases, operators find that it would have been useful if some of these parameters had alarms as well (Vicente & Burns, 1995). This suggests that the criteria that have been used for defining an alarm state are not ideal, probably being based on designers' intuitive notions of what states or events are important. This notion is consistent with our review of the literature. There is very little discussion of how alarm states have been, or should be, defined. One trend that was noted is that the advent of distributed control systems makes it easier for designers to define alarms, and thus, the total number of alarm states and events that have been specified seems to have increased as a result (Zwaga & Hoonhout, 1994). As Bray (1989) puts it, the attitude seems to be: "If in doubt, alarm it" (p. 212)!

The exception to this relative silence in the literature is the functional approach to alarm definition outlined by Goodstein (1985). This work builds on the abstraction hierarchy

framework developed by Rasmussen (1985), which provides a way to represent a plant at 5 levels of abstraction: Functional Purpose (the purposes that the plant is designed to achieve), Abstract Function (the mass, energy, and information topology of the plant), Generalized Function (the basic functions that have been designed into the plant), Physical Function (the equipment that is available to implement those functions), and Physical Form (the spatial location and appearance of plant equipment). Each level describes the plant in a different language. Adjacent levels are connected by a means-ends relation so that the level above any level specifies the ends that the function or component of interest can achieve, whereas the level below specifies the means that are available for implementing the function or component of interest. Because of this means-ends relationship, the abstraction hierarchy is explicitly designed to support goal-directed problem solving and information search. Following the means-ends relations in a top-down fashion allows operators to focus their attention on the subfunction or component that is relevant to the purpose or function of interest at the time, rather than having to consider all possible system elements. Thus, the abstraction hierarchy constrains search in a psychologically meaningful manner, thereby allowing operators to "see the forest through the trees".

It is precisely this feature which, Goodstein (1985) observed, makes the abstraction hierarchy useful for defining alarms. Rather than just defining context-free limits on individual parameters, as the single-sensor-single-alarm approach does, the functional approach to alarming provides alarms at higher levels of abstraction by integrating lower level data in a functional manner. For example, to determine whether a particular function is in an acceptable state requires that one examine and integrate data from a large number of components that are used to implement that particular function. In essence, this integration provides a form of context-sensitivity because the status of any one variable relevant to that function depends on the status of many, if not all, of the others. The functional approach to alarming takes advantage of the means-ends links in the abstraction hierarchy by integrating low level data into the functional units defined by those links. The result is a systematic approach to defining alarm states.

This approach has a number of other advantages as well (Goodstein, 1985). First, it results in a smaller number of alarms because of the context-sensitivity introduced by the integration process described above. Second, it provides a more meaningful set of indications because the alarms that are presented are explicitly defined with respect to plant purposes and functions. Third, it can help fault diagnosis as well because the links between levels of abstraction provide paths by which faults can propagate in a bottom-up fashion over time. Fourth, it also helps operators distinguish between common mode failures and multiple, independent faults. The former are revealed by multiple symptoms affecting a common function, whereas the latter are revealed by multiple symptoms distributed across functionally independent components (see Goodstein, 1985 for an example). Fifth, presenting alarms at multiple levels of abstraction also provides operators with an alarm overview since higher levels of abstraction have less detail than lower levels. Sixth, alarms at higher levels of abstraction are generally more important than those at lower levels (e.g., a threat to the plant purpose of safety is more important than the fact that an individual component is outside of its nominal range), so this approach automatically prioritizes alarms according to their functional importance. Finally, as the title of Goodstein's paper indicates, this approach also integrates alarms with information retrieval. Because the levels in the abstraction hierarchy are connected by means-ends links, operators will be able to follow these links to gather information that is pertinent to an active alarm. For example, if a certain function is being threatened, the operator will need to know where they should look next to get more detailed information about the components and systems that implement that particular function. With the functional approach, the operator need only consult the information at the next lowest level of abstraction that is connected to the alarmed function since only this information is pertinent to that function. In this way, the functional approach integrates alarm definition, information retrieval, and system status displays into a coherent framework.

Although the intended advantages of this approach are numerous, as far as we know, the ideas have not been empirically evaluated in any rigorous way on a representative scale. Nevertheless, this approach seems to have influenced the design of several advanced alarm

systems (see below), although none provide a faithful, systematic implementation of the philosophy. Although several of these advanced designs are based on hierarchical representations, none seem to include the level of Abstraction Function which describes the plant in terms of first principles. This is an important omission since first principles provide a reliable, generalizable basis for dealing with events which are unfamiliar to operators and which have not been anticipated by system designers.

Alarm Processing Techniques

A number of advanced alarm techniques have been developed in response to the deficiencies of traditional alarm systems. In this section, we will review a number of techniques that are intended to process alarm signals and information, independent of how the resulting output is displayed to operators. First, we will describe the methods themselves, and then we will describe the empirical evidence bearing on the utility of the various techniques.

Methods. One of the simplest alarm processing techniques that has been proposed to remedy one of the deficiencies of previous generation systems is that of lowpass filtering to eliminate the alarm "chattering" caused by a parameter oscillating in and out of its nominal operating range. This type of processing filters away threshold violations that occur within a very short period of each other, thereby removing one very annoying source of nuisance alarms in traditional CR designs.

Another alarm processing technique that is directed at reducing the number of nuisance alarms is based on logical filtering. In this case, the filtering is based on some type of logic, rather than the temporal properties of the signal itself as in the case of lowpass filtering. The logic provides a form of context-sensitivity to alarm processing. There are a number of different bases that can be used to define context. For instance, alarms can be filtered based on the plant mode (e.g., in a hot shutdown, a large number of symptoms that would normally be alarms at full power are expected and therefore normal). It is also possible to use other alarms as a context for filtering. For instance, if a hi-hi alarm on a particular parameter is activated, then the hi alarm is logically redundant and can therefore be filtered. Finally, it is also possible to use known

propagation paths or consequence alarms to filter alarms. As an example, during a reactor trip, a number of symptoms which would otherwise be abnormal are typically observed and are physical consequences of the trip, not independent contributing events. Note that this type of filtering is based on the known consequences of an event, not on operating mode.

A third advanced technique for alarm processing is the automatic prioritization of alarms. Again, this can be based on a number of different criteria. For instance, alarms could be prioritized according to their threat to safety. The functional alarming approach described earlier is one example of this type of prioritization. Another possibility is to prioritize alarms according to the time that the operator has available before he must respond to them. Although this may seem like a more pragmatically relevant criterion prioritization, it is not always possible to determine the time available for action. As with all other alarm processing techniques, prioritization schemes are only as effective as the assumptions on which they are based.

A fourth alarm processing technique that has been proposed to increase the informativeness of alarm systems is derivation of higher-order alarms from lower-level signals or alarms. There are at least two types of derivation algorithms, event-based and model-based (cf. Vicente & Tanabe, 1993). Event-based alarm derivation is predicated on the idea that one can identify a specific class of events by specifying a specific set of event pre-conditions. This can be implemented with a simple look-up table or with a more sophisticated knowledge-based programming techniques (see Kim, 1994 for a review). The automated disturbance analysis methods described earlier are an example. As we pointed out earlier, there are significant limitations associated with this class of techniques, the most important being that they do not provide the support that operators need to deal with unanticipated events (Kim, 1994).

Model-based alarm derivation is based, not on specific classes of events, but rather on a model which specifies certain relationships which must hold if the plant is operating normally (e.g., conservation of mass). As a result, model-based derivation techniques can be said to be event-independent (Vicente & Tanabe, 1993), and therefore are potentially relevant to a wide set

of disturbances (including unanticipated events), not just the finite set of event classes that have been explicitly identified by designers.

Empirical evaluations. Is there any empirical evidence to indicate that these novel alarm processing techniques lead to improved operator performance when compared to more traditional systems? Unfortunately, very few experiments have been conducted to evaluate these methods (at least under representative conditions), so the evidence available is quite meagre. For instance, we did not find any experiments evaluating the concept of lowpass filtering. The idea seems intuitively appealing, and could perhaps be implemented in a new alarm system without a great deal of engineering analysis. The potential obstacles to implementation seem to be the need for identifying the cutoff frequency for the various filters, and the computational load that the processing might impose.

As for logical filtering, an experiment was conducted at the Halden Reactor Project in Norway investigating several advanced alarm techniques, including logical filtering (Baker, Hollnagel, Marshall, & Øwre, 1985; Baker, Gertman, Hollnagel, Holmström, Marshall, & Øwre, 1985). The results indicated that performance with the alarm system with filtering was not significantly different from that with no filtering, despite the fact that the number of alarms during transients was reduced by approximately 50% with filtering. The authors note, however, that most of this filtering occurred later in the event, not at the beginning when operators initiate their fault management task.

With respect to alarm prioritization, Fujita (1989) conducted a controlled experiment comparing a traditional alarm system with an advanced alarm system that prioritized incoming alarms into one of three groups: normal status information, caution information, and alarm information. Note that all three categories of alarms were presented to operators (i.e., there was no filtering of alarms). Each of these groups was coded with a different colour (green, yellow, and red, respectively). The prioritization was performed according to a small number of relatively simple rules. These rules (in decreasing order of importance) are:

1). Handling by mode - alarms that are not relevant to the current mode are assigned to the normal status information group.

2). Handling by importance/magnitude - lower setpoint alarms for the same parameter (e.g., hi when hi-hi is activated) are assigned to the normal status information group. Similarly, less severe alarms for the same parameter (e.g., caution when failure is activated) are also assigned to the normal status information group.

3). Handling by cause-consequence - alarms that are a consequence of a component/system change that is already alarmed are assigned to the normal status information group. The exception is any case where the consequence alarm triggers an interlock action, in which case the alarm is assigned to the caution information group.

This is one of the more rigorously designed experiments conducted with professional operators and a full-scope simulator on alarm issues, so its results are of particular relevance.

Interestingly, the results revealed that there was no statistically significant difference between the two systems for identification of the initiating event. However, in 3 of 4 scenarios, the advanced alarm system led to significantly faster and less variable detection time of second malfunctions that occurred while operators were still busy dealing with the first fault. Also, the time it took operators to take control actions was also significantly improved with the advanced system for half of the scenarios. These results indicate that alarm prioritization can lead to improved performance (at least with the rules used by Fujita). They also suggest an important methodological conclusion, namely that the impact of alarm processing techniques may be mainly on subsequent malfunctions, not on initiating events. This is an important point to take into account in evaluating advanced alarm systems.

Finally, with respect to alarm derivation techniques, we were not able to find any experimental evidence. However, there is some development experience that is relevant. As mentioned earlier, efforts which have tried to derive an automated diagnosis of plant malfunctions based on processing of incoming signals and alarms have generally led to disappointing results (see Bray, 1989 and Kim, 1994). In contrast, efforts at using model-based

analysis to derive higher-order alarms have been promising. For instance, the early fault detection (EFD) system developed by Halden was implemented at the Imatran-Voima Loviisa plant in Finland and showed positive results in detecting plant failures at an early stage (Bye, Berg, & Øwre, 1994). These results demonstrate the feasibility of model-based derivation of alarms, but further empirical work is needed to confirm these preliminary findings.

Summary. A variety of alarm processing techniques have been suggested in an effort to overcome the limitations of traditional alarm systems. Unfortunately, very few experiments have been conducted to assess the value of these methods. The only positive conclusions we can make are that alarm prioritization can improve performance, and that model-based derivation of alarms seems to be a promising technique. Perhaps surprisingly, there is no evidence to indicate that alarm filtering improves operator performance. This does not mean that filtering may not be useful, but rather that it has not been shown to be so to date.

Alarm Presentation Techniques

In this section, we will review different methods that have been proposed or developed to present alarm information. Again, the methods themselves will be described first, followed by the empirical evidence bearing on the utility of the various techniques.

Methods. An obvious dimension of alarm presentation is the perceptual characteristics of the auditory cues that are selected to signal an alarm. The problem of developing guidelines for predicting appropriate loudness levels in complex noise environments seems to have been addressed (Patterson, 1982), and so will not be addressed further. However, some authors, most notably Gaver (1986), have suggested that the auditory channel can be used much more than it has been. Instead of just presenting a sound that indicates that something is wrong, more complex auditory stimuli can be developed to provide more information about the nature of the problem, and perhaps even where to look to get more detailed information for diagnosis and compensation.

Another dimension of the alarm presentation problem is whether lower priority or filtered alarms should be completely suppressed (and therefore not available to operators), or whether

those alarms should be made accessible to operators but in a less salient manner. Several alternatives for the latter approach would be to code them in a different colour (Fujita, 1989), have them available on request, or print them out on a printer so that they can be consulted if necessary.

A third presentation method is to integrate alarm information with process displays. This is made possible by the use of computer technology. Again, this can be achieved in several ways. For example, a parameter icon/label/value could flash on a CRT display to indicate that the value is abnormal. The same goal could be accomplished by changing the colour of the parameter icon/label/value to make it stand out from among the rest. A third possibility is to take advantage of the concept of emergent features (Bennett & Flach, 1992) to integrate the presentation of process status information and alarm information (see Vicente & Wang, 1996 for an example). With this approach, the distinction between an alarm and a status display essentially ceases to exist.

Another presentation technique that is intended to improve the informativeness of alarm systems is to organize alarms by function or system or task. The idea here is that collecting together alarms that belong to the same system or function or task should facilitate operator information search and retrieval compared to organizing the alarm information in some arbitrary way that is not as relevant to system purposes (Woods, 1991).

Finally, there is also the issue of whether alarm information should be presented in a parallel, spatially dedicated fashion that is continuously visible or in a more serial fashion that may or may not be spatially dedicated or continuously visible (O'Hara & Brown, in press). The former approach is that adopted in traditional tile-based annunciation systems, whereas computer-based alarm systems based on message lists have tended to adopt the latter approach. However, it is important to point that this choice is not necessarily tied to the implementation technology. For instance, it is possible to implement a parallel, quasi-dedicated alarm system with computer technology (e.g., Woods, Elm, & Easter, 1986; Easter & Lot, 1992). Rather than focus on the technology itself, it is more important to examine the implications of the particular

design for human performance. The potential advantages of the parallel approach are that operators can get an overview at a glance, can diagnose faults through pattern recognition, and know where to find any particular alarm (because it is always in the same place). The potential advantages of the serial approach are that it is more flexible so that alarms can be integrated with process displays, and grouped in various ways according to context (e.g., the task being performed). Of course, hybrid systems are also possible.

Empirical evaluations. As with alarm processing techniques, very few experiments have been conducted to evaluate the alarm presentation methods reviewed above under representative conditions. Consequently, there is very little basis upon which to determine which methods lead to improved performance.

Beginning with auditory cues, as far as we know, no experiments in the nuclear industry have been conducted to assess the value of using rich, multidimensional auditory information to indicate, not just that there is an alarm, but what the nature of the problem is. Gaver, Smith, and O'Shea (1991) conducted an evaluation with a process control microworld, and the results were very promising. Extension of these principles to the nuclear industry seems warranted.

As for suppression, the results obtained by Fujita (1989) in the evaluation mentioned earlier indicate that improved operator performance can be obtained without completely suppressing less important alarms. Furthermore, the studies by Kragt and Bonten (1983) and Vicente and Burns (1995, 1996) of operator use of alarms in practice suggest that completely removing alarms that are not considered to be urgent would cause severe problems. Operators rely heavily on such alarms as inputs to cognitive monitoring, and thus it seems inadvisable to remove this information in the absence of further research (cf. O'Hara, Wachtel, & Persensky, 1995).

There have been at least two studies which have investigated the effects of integrating alarms with plant status displays. The aforementioned Halden study conducted by Baker and colleagues (Baker, Hollnagel, Marshall, & Øwre, 1985; Baker, Gertman, Hollnagel, Holmström, Marshall, & Øwre, 1985) is one such study. Unfortunately, no statistically significant effects were observed, suggesting that there may be no benefits to be gained by integration. However, it

is possible that the problem is not with integration per se, but with the particular implementation. This possibility was investigated by a follow-up study conducted at Halden by Reiersen, Marshall, & Baker (1987). They compared a conventional annunciator tile system with 150 alarms, with a computer-based alarm system where alarms where presented in the context of an overview display on a graphic CRT display and in the context of more detailed, lower-level process displays. This advanced system was designed in such a way that the top level overview display and the lower level process displays were integrated, thereby aiding operators in selecting one of the latter by first examining the former. The results indicate that the advanced alarm system led to more efficient and more accurate information retrieval. However, the differences for fault detection speed were not as clear, with the advanced system being faster for 2 scenarios, the conventional system being faster for 1 scenario, and there being no difference on a fourth scenario. These results must be interpreted with caution since the conditions differed, not just in terms of alarms, but also in terms of process status displays as well.

Vicente and Wang (1996) have proposed a method for integrating alarms with process displays using emergent features. The logic behind this approach is identical to that used in model-based derivation (e.g., the early fault detection system developed by Bye et al., 1994). This form of integration has not been implemented on a large-scale system, although initial results with a process control microworld have been favourable. Given the feasibility of model-based derivation demonstrated at Loviisa, this approach to the integrated presentation of alarm and process information deserves further consideration.

With respect to organization of alarms, O'Hara and Brown (1991) review several studies which indicate that organizing alarms by system or by function has been shown to improve operator performance. In addition, operators have a strong preference for this organization scheme. This is not surprising since functional organization of displays is a basic principle of HF engineering.

The issue of serial vs. parallel presentation format is complex since it is actually a multidimensional issue (O'Hara, Wachtel, & Persensky, 1995). Kragt (1984) may have been the

first to investigate this issue experimentally, comparing a conventional tile-based alarm system with a more modern system consisting of a computer-based message list format. The study was conducted within the context of a simulated chemical process, which is considerably less complex than a nuclear power plant but much more complex than most laboratory studies. The results revealed that the computer-based alarm system was inferior to the conventional tile-based system in terms of various measures of fault management performance. The advantage of the more traditional design seems to be due to the fact that operators could interpret the information in the tiles at a glance, whereas with the alphanumeric message list they had to read the text messages to understand what state the plant was in. Furthermore, when alarms occurred rapidly in succession, the subjects with the message list were not able to retain an overview of the process whereas those with the tiles were. These results are important for two reasons. First, they show that "advanced" systems do not always lead to improved performance. Second, they show the value of a parallel, spatially dedicated presentation format in allowing operators to maintain an overview of the process and interpret alarms at a glance (cf. Woods, 1995).

According to O'Hara and Brown (1991), conventional tile-based displays have been found to be superior to CRT-based presentations during high-alarm density conditions in a number of other studies as well. Thus, the results obtained by Kragt (1984) seem to be generalizable.

Summary. A diverse set of alarm presentation techniques have been proposed. As was the case with alarm processing techniques, very few empirical studies have been conducted to assess the value of these techniques. The limited evidence available suggests the following conclusions:

- the possibility of using rich auditory information in alarms should be explored
- complete suppression of alarms that are not of highest priority is inadvisable
- the results on integration of alarms and process displays are equivocal, although future work in this area is warranted since information retrieval performance may be enhanced through integration
- alarms should be organized according to function or system

- alarms systems should include a parallel, spatially dedicated presentation format to support interpretation at a glance and maintenance of an overview of plant state.

Note that it is possible (probably desirable) to combine a parallel, spatially dedicated presentation format and integration of alarms with process displays into a single design.

New Advanced Alarm System Developments

A number of vendors worldwide have developed, or are in the process of developing, new alarm systems that incorporate one or more of the advanced alarm techniques described earlier. In addition to reviewing the literature describing these developments, O'Hara and Brown (in press) have observed these systems, and have had discussions with their designers. They describe 7 systems in total:

- ABB-CE NUPLEX 80+ Alarm System
- AECL CANDU Annunciation Message List Systems (CAMLS)
- EdF S3C Alarm System
- Halden Computerized Alarm System for HAMMLAB (CASH)
- Mitsubishi Dynamic Priorities Alarm System (DPAS)
- Toshiba and Hitachi ABWR Alarm System
- Westinghouse AWARE System

A description of each of these systems can be found in O'Hara and Brown (in press). We will limit ourselves to summarizing the new developments that characterize these systems as a whole.

The most obvious trend is towards integration of the alarm system with the remainder of the CR interface (e.g., overview panel displays, individual process displays). This integration is made possible by the move away from analog, hard-wired technology to digital, computer-based technology. It is important to note, however, that almost all of these new designs are hybrids in the sense that they consist of both traditional and advanced presentation media. Traditional tiles are usually used to provide an overview, whereas process displays and message lists on CRTs are used to provide more detailed information, thereby creating a hierarchical structure for the presentation of information. Another trend is towards the incorporation of advanced alarm

processing techniques (e.g., filtering, prioritization), despite the fact that the value of these techniques has yet to be clearly established empirically (see above). Although all of these systems are labeled "advanced", as far as we know, only one of them (the Mitsubishi design) has been empirically evaluated in a rigorous, representative manner with professional operators interacting with a full-scope simulator under a variety of challenging scenarios (Fujita, 1989). This is an important observation given the lack of industry experience with this type of technology.

CONCLUSION

The literature reviewed above clearly shows that many different alarm system design techniques have been proposed, but very little empirical research has been conducted. As a result, we do not have a very good understanding of the relative benefits and disadvantages of the techniques being proposed in terms of their impact on human performance and plant safety. Despite this, new alarm systems are being designed by nuclear power plant vendors world-wide. Therefore, there is an urgent need for research in this area. Furthermore, there is a great likelihood of having a significant impact on applied practice. In the meantime, the process criteria that have been proposed by the USNRC seem to be the most comprehensive and meaningful by which to evaluate proposals for future alarm system designs.

ACKNOWLEDGEMENTS

This work is a subset of the work performed for a research contract sponsored by the Atomic Energy Control Board of Canada. I would like to thank David Beattie (Humans*ystems*, co-principal investigator), Suzanne Rochford (contract monitor), Lawrence Lupton (AECL CRL), and especially John O'Hara (BNL) for their help.

BIBLIOGRAPHY

Baker, S., Hollnagel, E., Marshall, E., & Øwre, F. (1985). An experimental comparison of three computer-based alarm systems: design, procedure, and execution (HWR-134). Halden, Norway: OECD Halden Reactor Project.

Baker, S., Gertman, D., Hollnagel, E., Holmström, C., Marshall, E., & Øwre, F. (1985). An experimental comparison of three computer-based alarm systems: Results and conclusions (HWR-142). Halden, Norway: OECD Halden Reactor Project.

Bennett, K. B., & Flach, J. M. (1992). Graphical displays: Implications for divided attention, focused attention, and problem solving. Human Factors, 34, 513-533.

Bray, M. A. (1989). Alarm filtering and presentation. Nuclear Engineering and Design, 113, 211-218.

Bye, A., Berg, Ø., & Øwre, F. (1994). Operator support systems for status identification and alarm processing at the OECD Halden Reactor Project - Experiences and perspective for future development. In N. Stanton (Ed.), Human factors in alarm design (pp. 147-164). London: Taylor & Francis.

Easter, J. R., & Lot, L. (1992). Back-fitting a fully computerized alarm system into an operating Westinghouse PWR: A Progress report. In Proceedings of the IEEE Conference on Human Factors & Power Plants. Piscataway, NJ: IEEE.

Fujita, Y. (1989). Improved annunciator system for Japanese pressurized-water reactors. Nuclear Safety, 30, 209-221.

Gaver, W. W. (1986). Auditory icons: Using sound in computer interfaces. Human-Computer Interaction, 2, 167-177.

Gaver, W. W., Smith, R. B., & O'Shea, T. (1991). Effective sounds in complex systems: The Arkola simulation. In CHI '91 Conference Proceedings (pp. 85-90). Reading, MA: Addison-Wesley.

Goodstein, L. P. (1985). Functional alarming and information retrieval (Risø-M-2511). Roskilde, Denmark: Risø National Laboratory, Electronics Department.

Guerlain, S., & Bullemer, P. (in press). 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. Santa Monica, CA: HFES.

Kim, I. S. (1994). Computerized systems for on-line management of failures: A state-of-the-art discussion of alarm systems and diagnostic systems in the nuclear industry. Reliability Engineering and System Safety, 44, 279-295.

Kragt, H. (1984). A comparative simulation study of annunciator systems. Ergonomics, 27, 927-945.

Kragt, H., & Bonten, J. (1983). Evaluation of a conventional process-alarm system in a fertilizer plant. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, 586-600.

Lees, F. P. (1983). Process computer alarm and disturbance analysis: Review of the state of the art. Computers and Chemical Engineering, 7, 669-694.

Lupton, L. R., Lapointe, P. A., & Guo, K. Q. (1992). Survey of international developments in alarm processing and presentation techniques. Paper presented at International Symposium on Nuclear Power Plant Instrumentation & Control. Tokyo, Japan.

Marshall, E., & Baker, S. (1994). Alarms in nuclear power plant CRs: Current approaches and future design. In N. Stanton (Ed.), Human factors in alarm design (pp. 183-191). London: Taylor & Francis.

Mumaw, R. J., Roth, E. M., Vicente, K. J., & Burns, C. M. (1996). Cognitive contributions to operator monitoring during normal operations - Phase 2 (AECB Final Report). Pittsburgh, PA: Westinghouse Science & Technology Center.

O'Hara, J. M. (1994). Advanced human-system interface design review guideline: General evaluation model, technical development, and guideline description (NUREG/CR-5908, vol. 1). Washington, DC: USNRC.

O'Hara, J. M., & Brown, W. S. (1991). Nuclear power plant alarm systems: Problems and issues. In Proceedings of the Human Factors Society 35th Annual Meeting (pp. 1233 - 1237). Santa Monica, CA: HFS.

O'Hara, J. M., & Brown, W. S. (in press). Advanced alarm systems and human performance (BNL Tech. Rep. A3967). Upton, NY: Brookhaven National Laboratory, Department of Advanced Technology.

O'Hara, J. M., & Wachtel, J. (1991). Advanced CR evaluation: General approach and rationale. In Proceedings of the Human Factors Society 35th Annual Meeting (pp. 1243 - 1247). Santa Monica, CA: HFS.

O'Hara, J. M., Wachtel, J., & Persensky, J. (1995). Advanced alarm systems: Display and processing issues. In Proceedings of the Topical Meeting on Computer-Based Human Support Systems: Technology, Methods, and Future (pp. 160-167). La Grange Park, IL: ANS.

O'Hara, J. M., Brown, W. S., Higgins, J. C., & Stubler, W. F. (1994). Human factors engineering guidance for the review of advanced alarm systems (NUREG/CR-6105). Washington, DC: USNRC.

O'Hara, J. M., Brown, W. S., Stubler, W. F., Wachtel, J. A., Persensky, J. J. (1996). Human-system interface design review guideline (NUREG-0700, Rev. 1). Washington, DC: USNRC.

O'Hara, J. M., Brown, W. S., Baker, C. C., Welch, D. L., Granda, T. M., & Vingelis, P. J. (1994). Advanced human-system interface design review guideline: Evaluation procedures and guidelines for human factors engineering reviews (NUREG/CR-5908, vol. 2). Washington, DC: USNRC.

Patterson, R. D. (1982). Guidelines for auditory warning systems in civil aircraft (CAA paper 82017). London: Civil Aviation Authority.

Rasmussen, J. (1976). Outlines of a hybrid model of the process plant operator. In T. B. Sheridan and G. Johannsen (Eds.), Monitoring behavior and supervisory control (pp. 371-383). New York: Plenum.

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decision making and system management. IEEE Transactions on Systems, Man, and Cybernetics, SMC-15, 234-243.

Reiersen, C. S., Marshall, E. C., & Baker, S. M. (1987). A comparison of operator performance when using either an advanced computer-based alarm system or a conventional annunciator panel (HPR-331). Halden, Norway: OECD Halden Reactor Project.

Stanton, N. (1994a). Human factors in alarm design. London: Taylor & Francis.

Stanton, N. (1994b). A human factors approach. In N. Stanton (Ed.), Human factors in alarm design (pp. 1-10). London: Taylor & Francis.

Stanton, N. A., & Baber, C. (1995). Alarm-initiated activities: An analysis of alarm handling by operators using text-based alarm systems in supervisory control systems. Ergonomics, 38, 2414-2431.

Sultan, L. H., Lapointe, P. A., Guo, K. Q., & Lupton, L. R. (1990). Developments in alarm processing and presentation techniques (COG R&D Tech Report). Chalk River, ON: AECL Research.

Usher, D. M. (1994). The alarm matrix. In N. Stanton (Ed.), Human factors in alarm design (pp. 139-145). London: Taylor & Francis.

USNRC (1994). Human factors engineering program review model (NUREG-0711). Washington, DC: USNRC.

USNRC (1981). Guidelines for CR design reviews (NUREG-0700). Washington, DC: USNRC.

USNRC (1984). Standard review plan (NUREG-0800, Rev. 1). Washington, DC: USNRC.

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). Cognitive functioning of CR operators during normal plant operating conditions (AECB Final Report). Toronto: University of Toronto, Cognitive Engineering Laboratory.

Vicente, K. J., & Tanabe, F. (1993). Event-independent assessment of operator information requirements: Providing support for unanticipated events. In Proceedings of the American Nuclear Society Topical Meeting on Nuclear Plant Instrumentation, Control and Man-Machine Interface Technologies (pp. 389-393). La Grange Park, IL: ANS.

Vicente, K. J., & Wang, J. H. (1996). Taking full advantage of process constraints in advanced interface design. In Proceedings of the 1996 American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies (pp. 405-411). La Grange Park, IL: ANS.

Woods, D. D. (1991). The cognitive engineering of problem representations. In G. R. S. Weir and J. L. Alty (Eds.), Human-computer interaction and complex systems (pp. 169-188). London: Academic Press.

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

Woods, D. D., Elm, W. C., & Easter, J. R. (1986). The disturbance board concept for intelligent support of fault management tasks. In Proceedings of the International Topical Meeting on Advances in Human Factors in Nuclear Power Systems. LaGrange Park, IL: ANS.

Zwaga, H. J. G., & Hoonhout, H. C. M. (1994). Supervisory control behaviour and the implementation of alarms in process control. In N. Stanton (Ed.), Human factors in alarm design (pp. 119-134). London: Taylor & Francis.