Ergonomics, 38(11), 2371-2393, 1995

The Alarm Problem and Directed Attention in Dynamic Fault Management

David D. Woods
Cognitive Systems Engineering Laboratory
The Ohio State University
Columbus OH 43210 USA

Abstract

This paper uses results of field studies from multiple domains to explore the cognitive activities involved in dynamic fault management. Fault diagnosis has a different character in dynamic fault management situations as compared to troubleshooting a broken device which has been removed from service. In fault management there is some underlying process (an engineered or physiological process which will be referred to as the monitored process) whose state changes over time. Faults disturb the monitored process and diagnosis goes on in parallel with responses to maintain process integrity and to correct the underlying problem. These situations frequently involve time pressure, multiple interacting goals, high consequences of failure, and multiple interleaved tasks. Typical examples of fields of practice where dynamic fault management occurs include flight deck operations in commercial aviation, control of space systems, anesthetic management under surgery, and terrestrial process control. The point of departure is the “alarm problem” which is used to introduce an attentional view of alarm systems as tools for supporting dynamic fault management. The work is based on the concept of directed attention -- a cognitive function that inherently involves the coordination of multiple agents through the use of external media. Directed attention suggests several techniques for developing more effective alarm systems.

Keywords: Directed attention, joint reference, alarms, cognitive systems, fault diagnosis.

1. Introduction: The Alarm Problem

Fault management, the identification of and response to abnormal conditions, is a major component of the human's role as supervisory controller of dynamic systems. Evidence from disasters (Kemeny 1979; FCC 1991), from field studies (Cook et al. 1991; Moll van Charante al. 1993), from design reviews (Cooper 1977; Fink 1984), from experimental investigations (Kragt 1984; Kragt and Bonten 1983; Sorkin, Kantowitz and Kantowitz 1988), and from mathematical models (Sorkin and Woods 1985) all point out that operational personnel can have difficulties identifying, prioritizing and responding to abnormal conditions despite the presence of various types of alarm systems (both traditional annunciators and computerized variants) and diagnostic aids.

Several factors in alarm systems have been identified that contribute to difficulties in fault management (Lees 1984): nuisance alarms, ambiguous or underspecified alarm messages, alarm inflation, “alarms” that indicate system status rather than anomalies are just a few. The temporal dynamics of alarms are also relevant. The periods where alarms are densest are also likely to be those time periods of highest cognitive load and task criticality for practitioners. It is precisely during these periods of high workload that technological artifacts are supposed to provide assistance. But this is the time period where nuisance alarms will occur, where poorly designed alarms will distract and disrupt other tasks, where diagnostic search is most difficult and where there is the need to act to prioritize information processing tasks. Together these factors constitute what can be called the <u>alarm problem</u> (Woods, O'Brien and Hanes 1987).

The following episodes illustrate vividly for us some of the dimensions of the alarm problem.

> "The whole place just lit up. I mean, all the lights came on. So instead of being able to tell you what went wrong, the lights were absolutely no help at all."
> Comment by one space controller in mission control after the Apollo 12 spacecraft was struck by lightning (Murray and Cox 1990).

> "I would have liked to have thrown away the alarm panel. It wasn't giving us any useful information."
> Comment by one operator at the Three Mile Island nuclear power plant to the official inquiry following the TMI accident (Kemeny 1979).

> "When the alarm kept going off then we kept shutting it [the device] off [and on] and when the alarm would go off [again], we’d shut it off."
> “... so I just reset it [a device control] to a higher temperature. So I kinda fooled it [the alarm]...”
> Physicians explaining how they respond to a nuisance alarm on a computerized operating room device (Cook, Potter, Woods and McDonald 1991).

> The present lunar module weight and descent trajectory is such that this light will always come on prior to touchdown [on the moon]. This signal, it turns out, is connected to the master alarm--how about that! In other words, just at the most critical time in the most

> critical operation of a perfectly nominal lunar landing mission, the master alarm with all its lights, bells and whistles will go off. This sounds right lousy to me. ... If this is not fixed, I predict the first words uttered by the first astronaut to land on the moon will be "Gee whiz, that master alarm certainly startled me."
> Internal memo trying to modify a fuel warning alarm computation and how it was linked to the alarm handling system for the lunar landing craft (Murray and Cox 1990).

> A [computer] program alarm could be triggered by trivial problems that could be ignored altogether. Or it could be triggered by problems that called for an immediate abort [of the lunar landing]. How to decide which was which? It wasn't enough to memorize what the program alarm numbers stood for, because even within a single number the alarm might signify many different things. "We wrote ourselves little rules like 'If this alarm happens and it only happens once, don't worry about it. If it happens repeatedly, but other indicators are okay, don't worry about it.'" And of course, if some alarms happen even once, or if other alarms happen repeatedly and the other indicators are not okay, then they should get the LEM [lunar module] the hell out of there.
> Response to discovery of a set of computer alarms linked to the astronauts displays shortly before the Apollo 11 mission (Murray and Cox 1990).

> "1202." Astronaut announcing that an alarm buzzer and light had gone off and the code 1202 was indicated on the computer display.
> "What's a 1202?"
> "1202, what's that?"
> "12...1202 alarm."
> Mission control dialog as the LEM descended to the moon during Apollo 11 (Murray and Cox 1990).

> "I know exactly what it [an alarm] is--it's because the patient has been, hasn't taken enough breaths or--I'm not sure exactly why."
> Physician explaining one alarm on a computerized operating room device that commonly occurred at a particular stage of surgery (Cook et al. 1991).

The alarm problem has long been recognized as significant, but persists despite many different attempts to attack it. Some assaults have focused on the <u>perceptual functions</u> of alarms. Are the perceptual signals that mark abnormal conditions in a monitored process discriminable from the background and each other (e.g., Patterson 1990)? Another perceptual function of alarms is the capacity of the signal itself to attract people's attention (exogenous control of attention). This alerting potential refers to characteristics of the signal (e.g., sudden motion or sudden sound bursts against a quiescent background) to force, to a greater or lesser degree, an observer to orient to the signal itself. Apprehending these perceptual signals conveys the message that some event has occurred which may be of interest.

Other efforts have addressed the informativeness of alarms. Do semantic aspects of alarm messages inform observers about the kinds of abnormalities present or help discriminate between normal and abnormal states or between different kinds of abnormal conditions? There are several ways that alarms can be uninformative. Alarm signals and messages may be underspecified and ambiguous as in several of the above examples. Sometimes alarms occur frequently and seem so familiar that the observer misses the cases where something truly unusual is happening. In other words, the number of false alarms is high relative to the number of times the alarm signals that the abnormal condition it monitors for actually is present (Sorkin and Woods 1985). Another reason alarms may be uninformative is the context sensitivity problem (Doyle et al. 1989). Alarms often indicate the status of a parameter, subsystem or component. However, current status may or may not be actually abnormal without reference to additional data which specifies the relevant context (as in the case above from Apollo 11).

In this paper I will view the alarm problem from a slightly different vantage point where the alarm system is seen as an agent that attempts to direct the attention of the human observer. I explore characteristics of alarm systems and computer-based systems that affect practitioners' ability to shift attention to potentially interesting new events in data rich, high tempo and multi-task situations. Directed attention can help us understand why practitioners continue to use some traditional alarm systems in actual practice and abandon or underutilize the functionality of many new computerized systems which purport to address the alarm problem. The directed attention perspective integrates at least parts of the perceptual, informative, and cooperative aspects of alarms into a single cognitive function that I have labeled pre-attentive reference.

1.1 Joint Cognitive Systems

Some have attempted to solve the alarm problem through an automation finesse; that is, if we automate diagnosis, then the need for the cognitive task of alarm interpretation is eliminated. Ironically, despite the large efforts to develop more autonomous machine diagnosis capabilities, past field experience at introducing new diagnostic aiding systems shows that the benefits of new technology on fault management have tended to be derived more from information handling and visualization capabilities (e.g., Mitchell and Saisi 1987; Malin et al. 1991). Systems that only embody autonomous diagnostic capabilities without regard to the coupling to human partners in fault management have generally failed in the field (Roth et al. 1987 studies one case; also consider the failure of the disturbance analysis system developments in the nuclear industry in the 1980's).

In this paper, I will explore a joint cognitive system analysis of dynamic fault management. The concepts introduced here have evolved based on observations of practitioners doing fault management in evolving situations. The ideas arose in attempts to understand results from a series of studies of dynamic fault management and human-automation interaction in the four domains of flightdeck operations in commercial aviation, control of space systems, anesthetic management under surgery, and terrestrial process control. Investigations have included field observation of new support systems (Moll van Charante et al. 1993), simulation studies of practitioner cognitive activities in simulated emergencies (Woods et al. 1987), case studies of system development (Malin et al. 1991), and the design of new aiding strategies (Potter and Woods 1991).

There are two critical components to a joint human-machine cognitive system analysis. First, inherent in human-machine systems is cooperation across multiple agents, across both human and machine agents and across multiple people (Woods et al. 1994; Hutchins 1994). Second, the cognitive demands of dynamic fault management make control of attention a critical cognitive factor (Rabbitt 1984; Gopher 1991). This means that developing cognitive tools that support control of attention during evolving fault management situations is an important and overlooked aspect of developing new alarm and diagnostic systems. This paper will define several concepts about the control of attention during cooperative work in dynamic fault management.

2. Attentional View of Fault Management

Dynamic fault management is practiced in a cognitively noisy world, where very many stimuli are present which could be relevant to the problem solver (Woods 1994). There are both a large number of data channels and the signals on these channels usually are changing (i.e., the raw values are rarely constant even when the system is stable and normal). In this data rich environment, the human practitioner must track evolving situations loaded with unanticipated and potentially threatening events, but they can be barraged with signals that could be in principle informative or interesting. As a result, operators must build and maintain a coherent situation assessment in a changing environment where multiple factors are at work including one or more faults, operator interventions and automatic system responses. Attentional control, given multiple interleaved activities and the possibility of asynchronous and unplanned events, is a fundamental part of fault management.

Experts need to be able to manage several threads of activity in parallel, devoting enough attentional resources at the appropriate time in order to keep each on track. Resource saturation may be threatened, especially at high tempo periods. Experts may suspend even usually important tasks in order to perform tasks critical to cope with the system vulnerabilities most relevant in that context or stage. They assign less important tasks to subordinates, or even forgo some tasks completely in favor of an activity critical for the particular circumstances. Strategies for managing multiple activities in domain specific ways are even part of expert training (e.g., in aviation inder the label of crew resource management and in anesthesiology especially in new crisis management training).

Events in the process and in the interface (e.g., alarms) can serve as interrupts, prompts or reminders. Interrupt signals and interrupt handling are important when one functions in a cognitively crowded world. As data changes and new events are noted, how does one or should one modify their current task or cognitive resource priorities? Understanding action in the face of diverse, changing and highly uncertain situations depends critically on understanding attentional processes and the dynamic prioritization of tasks. A critical criterion for the design of the fault management systems is how they support practitioner attention focusing, attention switching and dynamic prioritization.

The critical point is that the challenge of fault management lies in sorting through an avalanche of raw data -- a data overload problem. This is in contrast to the view that the performance bottleneck is the difficulty of picking up subtle early indications of a fault against the background of a quiescent monitored process. While this may be the bottleneck in some cases, field studies of incidents and accidents in dynamic fault management emphasize the problem of shifting attention to potentially informative areas as many data values are changing.

Given the nature of human perceptual and attentional processes and the large amounts of raw data available, human monitors focus selectively on one of many possible objects, themes, activities with respect to the monitored process. Thus, shifting the focus of attention across the data field is a fundamental characteristic of fault management tasks. Shifting the focus of attention in this context does not refer to initial adoption of a focus from some neutral waiting state (Kahneman 1973). In fault management, one re-orients attentional focus to a newly relevant event on a different data channel or set of channels from a previous state where attention was focused on other data channels or on other cognitive activities (such as diagnostic search, response planning, communication to other agents). Dynamic fault management demands a facility with reorienting attention rapidly to new potentially relevant stimuli.

In an attention-based approach to fault management, alarm handling and diagnostic systems are seen as cognitive tools to support the human fault manager's or the fault management team's control of attention. This assumes, after Gopher (1991), that control of attention is a skillful activity that can be developed through training or supported (or undermined) by the design of representations of the monitored process.

2.1 Directed Attention Across Agents

From an attentional or cooperative point of view, alarms can be seen as messages from one agent, a first stage monitor, to another, a second stage supervisory agent who monitors multiple channels and whose cognitive and physical resources can be under moderate to severe workload constraints (Sorkin and Woods 1985). An alarm signal functions as a message from a first stage monitor intended to direct the attention of the supervisor to some particular area or topic or condition in the monitored process. In effect, the attention directing signal says, “there is something that I think that you will find interesting or important; I think you should look at this.” The attention directing signal functions as a kind of potential interrupt signal intended to influence or shift the receiver's focus of attention.

From the point of view of the supervisor receiving directed attention signals, again given that the supervisory agent has multiple demands competing for attention, he or she must evaluate the interrupt signal in parallel with ongoing activities and lines of reasoning (i.e., it is a dual or really a multiple task situation). The receiver must use some partial information about the attention directing signal and the condition that it refers to, in order to ‘decide’ whether or not to interrupt ongoing activities and lines of reasoning. Some attention directing signals should be ignored or deferred; similarly, some attention directing signals should re-direct attention. The quality of the control of attention is related to the skill with which one evaluates interrupt signals without disrupting ongoing lines of reasoning -- knowing when the attention directing event signals ‘important’ information and when the attention directing event can be safely ignored or deferred, given the current context. Overall, the situation can be expressed in a signal detection theory framework (e.g., Sorkin and Woods 1985) where one can err by excessive false shifts or excessive missed shifts. Thus, in principle, two parameters are needed to describe attentional control: a sensitivity parameter that captures variations in skill at control of attention and a criterion parameter that captures tradeoffs about the relative costs and benefits of under-shifts versus over-shifts of attention. Note that framing the problem in terms of signal detection points out that even very sensitive systems for control of attention will show errors of false shifts or missed shifts of attention.

An attention directing signal refers to an event or condition in the monitored process -- ‘look at this’. How the signal influences the processes involved in the control of the supervisor’s attention may depend strongly on information about (a) the referent event or condition that the first stage monitor thinks is worthy of attention and (b) the context in which this event or condition occurred. Specific information about why the monitor evaluated the event as worthy of an attentional shift also may be an important part of the attention directing process.

It is important to emphasize that directed attention is inherently both cognitive and cooperative. Directed attention is a kind of coordination across agents where one agent can perceive and direct the attentional focus of other agents to particular parts, conditions, or events in the monitored process. This kind of coordination involves several forms of joint reference. One aspect of joint reference involves “referring to something with the intent of directing another’s attention to it” in a mentally economically way (Bruner 1986, p. 63). Joint reference functions in the other direction as well, in that one agent can perceive where and to what the other is directing their attention to, without attention demanding explicit communication on the part of either agent. Hutchins’ analysis of the cognitive system involved in speed control during a descent in air transport aircraft illustrates this aspect of joint reference (Hutchins 1991). The physical activities associated with tasks carried out by one agent are inherently available for the other pilot to pick up without requiring explicit intent to communicate and without disrupting either’s ongoing lines of reasoning. Joint reference inherently involves the external representations of the monitored process by which agents assess the state of and interact with the monitored process. One refers to some part, condition, event or task in the referent world through some shared external representation of the monitored process. As a result, one agent can direct or can pick up another’s focus/activity on some part, condition, event or task in the referent world in mentally economical ways.

A study by Pavard et al. (1989) illustrates the dynamic of directed attention across agents under workload. They found that directing another’s attention depended on being able to see what the other agent is doing in order for one agent to be able to judge when another was interruptible. In other words, interruptibility is a joint function of the new message and the ongoing activity. This requires one agent being able to see the activity of the other in enough detail to characterize the state of the other’s activities --

what line of reasoning are they on? are they having trouble? does their activity conform to your expectations about what the other should be doing at this stage of the task? are they interruptible?

In the task investigated by Pavard et al., monitoring the other agent's activities to judge when the other agent was interruptible represented a new and mentally effortful task. For the agent trying to pass on information or hand off tasks, the burdens associated with monitoring or judging the other agent's interruptibility could interfere with their other tasks. This system, with direct visual access between team members, could function under moderate workload as long as the interrupting agent could monitor the other agents' activities. But the receiver had limited means to queue demands for attention or to partially evaluate the demands in order to decide how to allocate resources, to defer some tasks to lower tempo periods, to decide when to interrupt ongoing lines of reasoning and activity, or to signal that they are in fact interruptible (e.g., dealing with a non-demanding or low priority task). Under higher workload, the demands of coordination for the interrupting agent increased and interference between this and other task demands increased, leading to performance problems.

Pavard and his colleagues aided performance during high demand situations by devising a medium for interaction that was mentally economical for both parties. The design eliminated the need for the interrupting agent to monitor and judge the interruptibility of the receiver; the receiver was able to inspect a queue of demands for attention at their own initiative, as other task demands allowed, and in parallel with other activities. From the point of view of the receiver, the study showed that their control of their focus of attention was improved by the ability to queue demands for attention and to partially evaluate the different attentional demands in parallel with ongoing activities in order to decide when to switch their focus in a context sensitive way.

3. Preattentive Reference

Note the paradox at the heart of directed attention. Given that the supervisory agent is loaded by various other task related demands, how does one interpret information about the potential need to switch attentional focus without interrupting or interfering with the tasks or lines of reasoning already under attentional control. We can state this paradox in another way: how can one skillfully ignore a signal that should not shift attention within the current context, without first processing it -- in which case it hasn't been ignored.

This paradox has a parallel in perceptual organization (e.g., Kubovy and Pomerantz 1981). "Since the processes of focal attention cannot operate on the whole visual field simultaneously, they can come into play only after preliminary processes have already segregated the figural units involved ..." (Neisser 1976, p. 89). Neisser termed these processes preattentive. "An important part of the preattentive processes, therefore, is the segregation of detailed stimuli into bundles or segments that can be attended to or rejected as a whole (Broadbent 1977, p. 112). "These rules [preattentive processes] produce perceptual units that have a high probability of corresponding to distinct objects in the scene ... (Kahneman 1973, p. 68). These organizational processes are thought to operate without attentional demands; otherwise, we are left with a paradox: how can attentional mechanisms focus in on a part of the perceptual field if it first requires attention to structure the field into the objects or groups to be operated on.

It is important to see the function of preattentive processes in a cognitive system as more than a simple structuring of the perceptual field for attention. It is also part of the processes involved in orienting focal attention quickly to “interesting” parts of the perceptual field (Rabbitt 1984; Wolfe 1992). Preattentive processes are part of the coordination between orienting perceptual systems (i.e., the auditory system and peripheral vision) and focal perception and attention (e.g., foveal vision) in a changing environment where new events may require a shift in attentional focus at indeterminate times. Orienting perceptual systems are critical parts of the cognitive processes involved in noticing potentially interesting events and knowing where to look next (where to focus attention next) in natural perceptual fields (Folk et al. 1992).

To intuitively grasp the power of orienting perceptual functions, try this thought experiment (or better, actually do it!): put on goggles that block peripheral vision, allowing a view of only a few degrees of visual angle; now think of what it would be like to function and move about in your physical environment with this handicap. Perceptual scientists have tried this experimentally through a movable aperture that limits the observer’s view of a scene (e.g., Hochberg 1986). Although these experiments were done for other purposes theoretically, the difficulty in performing various visual tasks under these conditions is indicative of the power of the perceptual orienting mechanisms.

The orienting perceptual systems function to pick up changes or conditions that are potentially interesting. This perceptual ability plays a critical role in supporting how we know where to look next. Both visual search studies (e.g., Rabbitt 1984; Folk et al. 1992) and reading comprehension studies (e.g., the review in Bower and Morrow 1990) show that people are highly skilled at directing attention to aspects of the perceptual field or the text being read that are of high potential relevance given the properties of the data field and the expectations and interests of the observer.

I have claimed that designers of a computer based information system are creating a kind of virtual perceptual field (Woods 1995). In typical systems, the proportion of the virtual perceptual field that can be seen at the same time (physically in parallel) is very very small. In other words, the viewport size (the windows / VDUs available) is very small relative to the large size of the artificial data space or number of data displays that potentially could be examined. This property is often referred to as the keyhole effect. Given this property, shifting one's “gaze” within the virtual perceptual field is carried out by selecting another part of the artificial data space and moving it into the limited viewport.

In these kinds of computer-based information systems, the designer has created a virtual perceptual field where the observer must function without the assistance of the orienting perceptual systems. When dealing with the computer medium, the burden is on the designer to explicitly build in mechanisms to support the operation of the orienting perceptual systems and the functions that they perform in a fully articulated cognitive system adapted to a changing environment. The design of the alarm system is a critical part of supporting the function of these mechanisms, and directed attention is one process that contributes to these functions. Preattentive processing is part of how we are able to achieve a "balance between the rigidity necessary to ensure that potentially important environmental events do not go unprocessed and the flexibility to adapt to changing behavioral goals and circumstances" (Folk et al. 1992, 1043). The concept of directed attention suggests that characteristics of external artifacts can influence the expression of this basic cognitive competency.

Now let us return to the paradox that we started with: how can one decide whether or not to ignore a signal without first processing it; or how can one decide whether or not a new signal warrants interrupting the ongoing line of reasoning without interrupting or interfering with that very line of reasoning. Skillful control of attention, knowing when the attention directing event signals 'important' information and when the attention directing event can be safely ignored or deferred given the current context, depends on the supervisory controller in an event-driven environment somehow being able to notice potentially interesting changes without drawing on or interfering with limited attentional resources. Analogous to preattentive processes in perception, some kind of preattentive evaluation of an attention directing signal is required to resolve the paradox. I claim that the ability to carry out a preattentive evaluation of attention directing signals depends in part on the characteristics of the alarm system and the representation of the monitored process. Thus, I will refer to the characteristics of alarm and display systems that support a preattentive evaluation of attention directing signals as attributes of preattentive reference.

Note that preattentive reference is a joint cognitive system property. It refers to more than cognitive processes within a single head. Preattentive reference is about how characteristics of the alarm system and characteristics of the external representations of the monitored process available to the supervisory controller support skilled control of attention in data rich, dynamic situations.

The criteria for preattentive reference are that the attention directing signals

(a) are capable of being picked up by the supervisory controller in parallel with ongoing lines of reasoning and ongoing activities,

(b) includes partial information on what the attention directing signal is referring to, so that the observer can pick up whether the interrupt signal warrants a shift in attention or not,

(c) the assessment of the interrupt signal and the partial information that it carries must be mentally economical and not require an act of focal attention.

In effect, these criteria demand the potential for 'peripheral access' in the sense of being able to evaluate the interrupt without interrupting ongoing lines of reasoning (not necessarily in the sense of literal detection via peripheral vision).

An example of a system meeting these criteria serendipitously occurred in power generation process control. The position of a device (rod position in the core of a nuclear reactor) under the control of an automatic system was indicated via mechanical counters. These counters happened to make an audible click sound when the mechanism changed values as the device position changed, that is, click rate varied as rods were moved into and out of the core to regulate and balance the heat generated (the same serendipitous auditory indications occurred with another system under automatic control which regulating the boron concentration in the coolant fluid).

Operators are able to monitor the auditory indications preattentively: (a) the indications were available in parallel with other indications (primarily visual), (b) the clicking pattern contained information about the activity of the system (is the system active or quiescent? making small or large adjustments?) that could be used to recognize whether or not system activity matched practitioners' expectations for the context, (c) practitioners could monitor the auditory indications in parallel with other activities and switch attention only when they recognized anomalous behavior in that portion of the monitored process.

Another example, again non-visual, is the role of the voice loops at the mission control center that manages space missions (e.g., Murray and Cox's 1990, description of the Apollo missions illustrates how the voice loops function as a coordinative tool). Outsiders see the voice loops as sheer cacophony; but space flight controllers see it as an essential tool. The voice loop system enables preattentive reference -- one can notice potentially interesting events or signals without drawing on limited attentional resources. Similarly, shared voice loops support preattentive reference in the context of aircraft carrier flight operations (Rochlin et al. 1987).

> "... everyone involved ... is part of a constant loop of conversation and verification taking place over several different channels at once. At first little of this chatter seems coherent, let alone substantive, to the outside observer. ... one discovers that seasoned personnel do not 'listen' so much as monitor for deviations, reacting to almost anything that does not fit their expectations ..." (Rochlin et al. 1987, p. 85).

Experience with the computerization of traditional annunciator type alarm systems provides us with a visual example of the properties of alarm systems and computer based information systems that contribute to preattentive reference. Annunciator systems consist of a number of backlit tiles. Each tile has engraved on it a message concerning some state or event in the monitored process. Each individual state change has a dedicated tile which associated with it, and the tiles are laid out in a fixed spatial array. When a state change or event occurs, the associated tile is backlit (see Woods et al. 1987 for examples of these kinds of systems). Field experience indicates that simple computerization of existing alarm systems has failed to correct the alarm problem and in some cases has exacerbated it (Woods et al. 1987; Potter and Woods 1991). I mean simple computerization in the sense of moving the exact same alarm messaging constructs from a physically parallel and spatially dedicated medium (backlit annunciator tiles each dedicated to a particular alarm triggering condition) to serial presentation in the computer medium (usually including chronologically ordered message lists). Potter and Woods (1991) describe at some length the deficiencies of chronologically ordered alarm message lists, so I will not dwell on them here.

While annunciator type alarm systems have a variety of deficiencies, they have some properties that seem to provide some practitioners, with large amounts of experience, with the potential to develop strategies for control of attention. Spatially dedicating each alarm message within a physically parallel space at least provides a minimal level of support for preattentive reference. When a new alarm occurs, it always occurs at some specific place within the spatial array of annunciator tiles. The spatial array provides some of the conditions of a natural perceptual field so that orienting perceptual mechanisms (visual and auditory) can pick up the fact that an event has occurred in the monitored process. The fact that each alarm state is dedicated to a single physical location (and, ideally, that functionally or structurally similar alarms tend to occur in neighboring locations) means that observers can know something about the event that occurred without invoking focal attention (in the case of nuclear power plants until the mid- to late 1980s, the engraving on the tiles were often not legible at the distances operators were stationed in the control room, yet operators could pick out patterns of events). In other words, it is possible for observers using these annunciator type systems to partially evaluate the attention directing signal to 'decide' whether a shift in attention is warranted or not.

Another study (Leroux in press) illustrates the paradoxical nature of preattentive reference. In this study air traffic controller teams directed aircraft in a busy sector with multiple conflicts between aircraft to resolve (the study was run on a full scope simulation facility with actual controllers). One aircraft in the sector behaved in an abnormal fashion (a very abnormally low rate of ascent). In one of the problems this aircraft's behavior was directly linked to the conflicts that the controllers needed to resolve; in another scenario the same aircraft and behavior occurred in a lower density part of the sector and did not interact with the conflict resolution tasks of the controllers. All of the teams noticed the aircraft's abnormal behavior when it interacted with their conflict resolution tasks; virtually all of the teams failed to notice the aircraft's abnormal behavior when it was not part of their conflict resolution tasks. Several teams insisted, when prompted in general and later in particular, that this aircraft had exhibited no extraordinary behavior. One crew, when watching a replay of the scenario, even insisted that the investigator had modified the playback -- they could not conceive of having missed the abnormal behavior. The same locally anomalous behavior was attended to (and therefore was reportable) when it is a part of the practitioner's larger goals and task context, and it was not attended to (and therefore was not reportable) when it was not a part of the practitioner's larger task context.

This case points out how preattentive reference is not a conscious decision or judgment but rather is some kind of recognition driven process. The observers are not aware of the normal or expected or irrelevant parts of the flow of activity, but they are capable of recognizing the anomalous when it is relevant to the larger context and their goals. They only see or pickup or are sensitive to the relevant portions of the data field, even though what is relevant varies with both properties of the data field and with the interests and expectations of the observer. Interestingly, the same conclusion can be drawn about visual search (Rabbitt 1984; Folk et al. 1992) and text comprehension (Bowers and Morrow 1990). While the concept that experts see problems in terms of meaningful high level structures is old (DeGroot 1965; Chase and Simon 1973), the concepts of directed attention and preattentive reference show how this ability is modulated by characteristics of the cognitive artifacts embedded within a joint cognitive system.

None of these examples are meant to indicate that the solution to the alarm problem simply is to mimic properties of traditional media for handling alarms such as annunciator panels or voice loops. Instead, the point is to understand how properties of these media interact with the cognitive demands of and strategies for dynamic fault management. It is this deeper understanding that can support skillful rather than clumsy use of technological possibilities to create more effective joint cognitive systems.

Note that my characterization of directed attention and preattentive reference is not an explanation of the internal cognitive mechanisms that produce it, but rather a characterization of how it functions in the larger context of practitioners handling dynamic fault management problems with various kinds of cognitive artifacts. Artifacts that support preattentive reference could be mentally economical in several ways: processing these displays could be literally preattentive; the artifact could simply support timesharing (task switching) in a way that minimizes disruption of the ongoing lines of reasoning and activities; the artifact could support parallel processing using different resources (e.g., an ongoing task that primarily uses the visual channel with interrupt signals occurring on auditory channels). Whichever mechanism or combination of mechanisms underlies preattentive reference in general or in a given case, the key component is the way that artifact helps orienting perceptual systems coordinate with focal attention in natural perceptual fields where new events worthy of attention can occur at indeterminate times.

4. Attention Directing Displays and Preattentive Reference

4.1 Underspecified Alarms

In general, an attention directing signal says, “there is something that I think that you will find interesting or important that you should look at.” Let us assume that the perceptual properties of the signals are well designed so that observers can apprehend that a signal is present relative to the perceptual background and yet not to be so powerful as to disrupt all processing temporarily (as in the startle response to very loud and sudden sounds). There are essentially two parts to the content behind an attention directing message. First, there is the state, event, or behavior of the monitored process that is being referred to -- what is it that the monitor thinks is worthy of attention. Second, why does the alerting monitor think that this condition or event is interesting. This is almost always based on a difference relative to a background or contrast state of the monitored process: either a contrast between actual and desired state of the monitored process (an abnormal condition) or a contrast between actual behavior and the agent's model or expectations about current or future process behavior (an unexpected condition).

In addition, an abnormal or unexpected event varies in importance depending upon its relationship to the larger context of goals and competing activities. For example, an alarm may refer to a valve position -- valve X is closed. But if this all that is noted in the attention directing signal one must infer on their own what is anomalous about this valve position. It could be anomalous because the system of which it is a part could not perform its function if it were needed; in other words, the redundant backup mechanism is now unavailable. Or it could mean that the system of which it is a part is no longer performing its function; in other words, the process for meeting some constraint on the monitored process has malfunctioned. These are two different kinds of anomalies in the process; they both refer to the same state of this component, but in different larger contexts.

Problems occur when alarms messages say only that "there is trouble, here" without any further indication of the specific anomaly detected or the function that is impaired ('here' refers to some aggregated unit of description of the monitored process, typically traditional system / subsystem divisions). These kinds of alarms are underspecified and provide no support for control of attention. The observer has to shift attention away from the current line of reasoning and search for more information to clarify what kind of anomaly is being referred to and to evaluate its importance in the current context. The alternative to this break is to defer evaluation of the signal to a lower tempo moment.

One common example of this dynamic occurs when designers of computer-based information systems mark a digital value or component icon yellow or red when the machine monitor detects an anomaly in the monitored process, without providing any further information concerning the nature of the anomaly (Malin et al. 1991). It appears to be a virtually universal folk concept that changes in hue alone are an effective directed attention signal. While a hue change could be an effective perceptual cue to signal that an event has occurred, the hue change alone provides no indication whatsoever of the kind of trouble that triggered the attention directing signal. The hue change is in effect a group alarm or an aggregated attention marker, i.e., many different kinds of anomalies could be the source of the signal. One cannot evaluate a group alarm, in these kinds of designs, without interrupting ongoing lines of reasoning and devoting focal attention (a) to determine where more data about the anomaly resides in the virtual perceptual field hidden behind the narrow keyhole of the VDU, (b) to maneuver to that display, and (c) to evaluate the data relative to the current context (Woods 1995). Thus, using hue as a group alarm fails to support directed attention when multiple signals and events are occurring which may indicate the need for the observer to shift attention. The group alarm does not provide a partial, peripherally accessible indication of the shape or type of information that the attention director thinks is interesting or why it thinks it is interesting.

An investigation of the impact of a computer display that consisted only of icons representing aggregations of components which changed hue when the machine monitor detected any anomaly within that aggregation of components revealed the weakness of this technique (Reiersen, Marshall and Baker 1988). One of the characteristics of dynamic fault management is the cascade of disturbances that follow from a fault which can produce an even greater cascade of low level messages about potential anomalies (Woods 1994). As a result, group alarms hide most of the changes going on in the monitored process that result from the disturbances spreading from one or more faults. Reiersen et al. found that operators of a simulated nuclear power plant did not utilize the hue coded icon display preferring other displays that provided greater depth of data about the state of the plant in a single view. The hue coded icon display provided very little data, forcing the operator to switch to other displays as soon as any trouble at all occurred in the monitored process; in other words, it was a data sparse display. Field studies support this result. Practitioners treat systems with uninformative alarm systems as if there were only a single master caution alarm (Moll van Charante et al. 1993).

This case illustrates that attention directing signals should refer to the specific anomaly that is the basis for the monitor issuing an alert. In addition, it could include or should at least make it easy, both physically and cognitively, for the practitioner to see the relationship of this anomaly to the larger context. Furthermore, it is very important not to confound the alerting signal itself (the look or orient part of the message) with indications about what is the potentially interesting condition or event.

The above ideas were utilized in the development of one system for supporting alarm management (Woods et al. 1986). This system has attention directing displays which signal the kind of anomaly indicated as well as the specific anomalous event or condition. The alarm system consisted of an array of computerized alarm 'tiles' for each function of a particular monitored process. Each function-based tile signaled the kind of anomaly or disturbance detected by the machine monitor as well as the specific evidence on which this was based. In this system three classes of anomalies were defined: (a) constraint or goal violations (e.g., constraints on providing sufficient material inventory in a reservoir), (b) process disturbances including both the absence of a desired influence (e.g., the failure of an automatic system to respond as needed) and the presence of an undesired influence (e.g., a piping break), and (c) process unavailabilities (backup systems that would not function if needed, for example, loss of a vital support system).

The soft tile for each function occurred in one fixed position and each tile was subdivided into three fixed areas, one for each of the anomaly classes (spatial dedication). When a specific triggering event occurred, the observer received several indications in parallel: the overall soft tile changed to indicate an event in that portion of the monitored process; the relevant portion of the tile dedicated to the kind of anomaly detected in each case changed; and a computer-based message occurred in that portion of the soft tile which specified the condition detected by the machine monitor. Furthermore, observers could see relationships between functions and how they changed over time because the soft tiles formed a spatial array. These characteristics were developed consciously to support preattentive reference. Observers could in principle extract partial information about the kind of event and its importance in the larger context without disrupting ongoing lines of reasoning.

4.2 Nuisance Alarms

A frequently noted contributor to the alarm problem is nuisance alarms. From a cognitive systems point of view, a nuisance alarm is a signal that attempts to direct attention to an event, but that event is frequently one that does not warrant a shift in the agent's focus of attention. In other words, a nuisance alarm is a consistently false signal for an attention shift. For example, during the descent phase of flight on a B-727 aircraft the flight engineer will know when to grab the alarm silence control to anticipate and prevent a nuisance alarm from sounding that is triggered in every descent.

It is important to distinguish between an attention directing signal that turns out to not warrant an attention shift in this particular case (but could have) and an attention directing signal that does not warrant attention shifts consistently in a specific context. The latter is a nuisance alarm in that context. Nuisance alarms represent a cost to practitioners because they frequently constitute an interruption of attention to other tasks and lines of reasoning. Frequent nuisance alarms is a cue that suggests sharpening the definition of what is potentially interesting for this context -- an increase in the "intelligence" or informativeness of the alarm system.

The nuisance value of alarms also depends on the context of practitioner activity in which the alarm occurs. Many fields of practice involve busy periods where practitioners cannot devote their attention exclusively to one subsystem or condition. In one study (Moll van Charante et al. 1993) the pace of the activity was faster than the time required for the first stage monitor to detect and announce faults. This meant that practitioners were doing other things by the time an alarm appeared so that alarms always constituted an interruption of practitioner attention to other tasks. Users would initiate or intervene on one task and then proceed to another task. Later they would be alerted by an alarm that the previous action had been unsuccessful. This alarm necessitated stopping ongoing activity, switching attention, locating the appropriate subsystem, recall of the prior actions, and troubleshooting of the subsystem even when state or event signaled was of low priority relative to other events.

4.3 Context Sensitivity

The same attention directing signal can and sometimes should evoke different responses in different contexts. The meaning of the signal (and the underlying anomaly) depends on what other problems are present in the monitored process and on where the practitioners are in the process of assessing and responding to the disturbances.

Typically, alarm system designers attempt to help the practitioner know when an anomaly should interrupt ongoing activity by developing a fixed, static priority assignment to individual alarm signals. Usually, two or three classes of priority are defined and then individual alarm signals are assigned to one these categories. Presumably, there are only a few high priority alarms that occur in the same time period and alarms in the lower priority classes do not need to be processed in order to evaluate the significance of the high priority ones. In other words, the static priority technique tries to cope with alarm handling demands through a scale reduction process.

But scale reduction does not directly aid directed attention. The context sensitivity problem makes it very difficult to define one static set of priorities. Alarms should help link a specific anomaly into the larger context of the current activities and goals of supervisory agents. What is interesting depends on practitioners' line of reasoning and the stage of the problem solving process for handling evolving incidents. Examples of how the supervisor's situation assessment or mindset affects the interpretation of an alarm include:

- if the background situation assessment is 'normal system function,' then the alarm is informative, in part, because it signals that conditions may be moving into a qualitatively different phase of operations -- abnormal or emergency operations;
- if the background line of reasoning is 'trying to diagnose an unexpected finding,' then the alarm may be informative because it supports or contra-indicates one or more hypotheses under consideration;
- if the background line of reasoning is 'trying to diagnose an unexpected finding,' then the alarm may be informative because it functions as a cue to generate more candidate hypotheses that might explain the anomalous process behavior;

- if the background line of reasoning is 'executing an action plan based on a diagnosis,' then the alarm may be informative because it functions as a cue that the current working hypothesis may be wrong or incomplete since the monitored process is not responding to the interventions as would be expected based on the current working hypothesis.

In other words, the context sensitivity of interrupts is the major challenge to be met for the development of effective alarm systems, just as context sensitivity is the major challenge for developing solutions that treat any data overload problem (Doyle et al. 1992; Woods 1995).

4.4 How Can Attention Directing Displays Support Preattentive Reference?

The claim here is that directed attention, as a part of the skilled control of attention, can be supported or undermined by characteristics of the interface between practitioner and monitored process. Understanding directed attention should point towards techniques or principles for the design of attention directing systems and displays. How can attention directing displays support partial, preattentive characterization of potential interrupts? The key is to think about the way that displays, representations, alarms, etc., can engage the orienting perceptual systems which coordinate with focal attention.

For the case of peripheral vision, my observations of directed attention in field settings suggest several properties of dynamic representations of the monitored process that are likely to support partial, preattentive characterization of potential interrupts. A first requirement is that the representations of the monitored process must capture and emphasize change and events (Woods 1995), given that peripheral vision is relatively good at picking up change and given that the unit of analysis for directed attention is knowing which changes in the process are interesting or important to focus attention on. Second, representations of the monitored process that will support preattentive reference need take advantage of nonvisual channels or, if visual, support check reading or peripheral visual cues. Third, analog pattern oriented graphic representations are likely to better support peripheral access (Sorkin et al. 1988) -- in part, the patterns generally will be lower spatial frequency stimuli and the processing required to pick up patterns may be mentally economical in one or more of the ways outlined earlier.

4.4.1 *Spatial Dedication:* A powerful technique is spatial dedication, i.e., data on the status of some part of the monitored process appears in a fixed and consistent location. Thus, if one picks up a change in that location, then they are aware implicitly that an event has occurred on the relevant subsystem, function or data channel. In three different studies of the introduction of multi-function computer based display systems for highly dynamic and event driven fields of practice (a medical operating room case, a space flight control case, and a study of experienced users of extended spreadsheets), it was found that practitioners made minimal use of the flexible display of data afforded by the computer medium and instead converted the displays into a fixed and spatially dedicated form where the same data and topics always occurred in the same places.

Note that spatial dedication is an implicit property of hardwired media for representation where physical constraints require each display to occupy one fixed position in a larger field. But the computer medium, as typically used, eliminates spatial dedication as an organizing cue. Given (a) the typical kinds of graphic forms designed for computers (which do not highlight change and events), (b) the fact that the basic unit of data display remains the raw value, (c) the keyhole property of the computer medium, and (d) the absence of spatial dedication, the result of this combination is little or no support for preattentive reference. This is not to imply that one should simply return to spatially distributed and dedicated annunciator systems to support directed attention but rather to note how such systems provide some support for this cognitive function. Understanding the partial success and the limits of these systems can lead us to new ideas for how to create preattentive reference and support directed attention when the substrate is the computer medium.

4.4.2 *Auditory Displays:* We can think about how representations using the auditory channel can support preattentive reference by considering examples like the serendipitous auditory indications that reflected the behavior of automatic control systems or the function of the voice loops in space system control centers. In one sense, auditory indications can support preattentive reference simply because vision is often the dominant channel for monitoring system status. More fundamentally, auditory indications may be especially relevant in building an implicit awareness of orientation in an environment (though we wish to apply it to an abstract environment of the changing status of a dynamic process). Audition is special in that it is an inherently temporal modality; thus, it may have some special advantages for representing the behavior of a process (behavior as a change over time). One approach is to link auditory stimuli to the behavior of a system or process (Gaver 1991). If audible consequences result from the behavior of the process (e.g., clicks that occur when a system changes state) and if the auditory indications correspond to meaningful categories of process behavior (e.g., is the system active or quiescent? is the trend deteriorating or recovering? is system state changing in one direction or another?), then the auditory representation has provided the basis for practitioners to monitor the system preattentively.

4.4.3 *Analog alarm displays:* Alarm signals are usually categorical -- a condition is present, the value has crossed a discrete threshold. One way to support preattentive reference may be to make greater use of analog representations with visual or auditory highlighting of discrete category shifts (see Woods et al. 1987 and Sorkin et al. 1988 for examples). By showing continuous values with respect to an underlying dimension, observers may be able to better recognize when an attention directing signal warrants interruption of ongoing activity. For example, practitioners may react differently to a rapidly deteriorating situation as compared to one that is stable but abnormal, depending on the larger context.

4.4.4 *The 'Intelligence' of Alarms:* In the directed attention paradigm increasing the 'intelligence' or informativeness of an alarm system is sharpening the ability of the first stage machine monitor to identify potentially interesting conditions/events as a function of context (see Doyle et al. 1989; Doyle et al. 1990 and Doyle et al. 1992 for one line of research to develop such systems). Informativeness lies in part in the first stage monitor's ability to answer the question -- when are events in the monitored process worthy of attentional shifts? In other words, attention directing signals are ambiguous when it is difficult to know on the basis of the signal alone whether an attentional shift is really warranted.

But remember that in dynamic processes variation and change is the norm. Which of these variations or the absence of variation mark an interesting event? The most basic factor that governs when new events are interesting is a difference against a background, including both departures from normal function and departures from models of expected behavior in particular contexts. Recognizing such mismatches can be very difficult as when the interesting behavior is the absence of an expected change or when one change initiates a set of changes, for example, one event can shift the reference state against which many other states should be evaluated.

Let us see these points concretely in an example taken from the development of an intelligent system for monitoring satellite communications. The intelligent system detects and issues a warning (a message appears in the alarm window and an icon changes to red to indicate a problem in transmission) when it detects a loss of signal condition (Figure 1, panels A and B). Obviously, a loss of transmission from the satellite is an abnormal condition -- or is it always? Space flight controllers complained that the intelligent system was emitting a nuisance alarm under some conditions. The intelligent system thought there was an abnormal loss of signal and tried to diagnose the underlying fault when, in fact, it was normal and expected to lose satellite transmissions during a transponder switchover, a normal, expected and scheduled event. Figure 2 graphically illustrates the sequence of events. The initial response of the knowledge engineers was to adjust the knowledge base to eliminate all messages and graphical indications of the occurrence of the event when it was not abnormal. Now the representation provided no indication of the transponder switchover event, and no indication that the expected loss of signal had occurred (the display looked like Figure 1, panel A during the switchover). However, users complained about this approach as well because there was now no indication at all that the event was occurring as expected. Not only was there now no basis for seeing the event, there was no way to see departures from the expected pattern, for example, was the expected resumption of signal transmission delayed or did the second transponder pick up as expected? Ideally, the solution was to animate a representation like that illustrated in Figure 2 which would provide a view of the context (scheduled transponder switchover) and the expected pattern so that the machine monitor could direct the human controller's attention to abnormalities or unexpected behavior against that common frame of reference. Note that this approach to enhancing the intelligence of alarm systems depends on different kinds of representations of the monitored process that provide enhanced visualization of function and malfunction (Woods 1995) and that this approach is very different from the approach of automating fault diagnosis (Malin et al. 1991).

4.4.5 *Joint Reference:* Directed attention is in part about the coordination across agents. This kind of coordination involves joint reference -- referring to something with the intent of directing another's attention to it. The attention directing signal simply could indicate something interesting or important has happened (e.g., a master warning alarm). It could indicate that something interesting is going on here (e.g., look here). In these forms of joint reference the first stage monitor is saying in effect, 'something that I think that you will find important has occurred, ask me when you get a chance and I'll tell you what is interesting and maybe even why.' However, the attention directing signal can refer in part to the state or behavior that is being labeled 'interesting or worthy of attention' (what is interesting). And it could include reference to, or occur in a larger representational context that shows, the basis for why this behavior is interesting in this context.

Note that joint reference implies some kind of shared representation of the monitored process. What are the requirements for a shared representation of the monitored process that supports joint reference? First, it requires an externally stored representation of conditions and events in the referent world available to all the agents (open tools -- Roth et al. 1987; Hutchins 1991) . This shared external representation must support an economical description or depiction of what aspect or behavior is being referred to as opposed to having the agents explicitly describe to one another what is being referred to and what is interesting about the thing-referred-to. The shared representation must also support an agent’s ability to refer to something that is visible (perceivable) to the other agent and that is unambiguous as to what is being referred to. Note that a shared external representation also assumes that there is a shared mindset across the cooperating agents about the background field against which the agents can all recognize interesting conditions or behaviors.

A shared external representation of the referent process is more than a display that is physically available to all of the agents. It also refers to what can be seen about the referent process and to the ability to support joint reference. For example, on computerized flightdecks pilots interact with the automation through a control-display unit (CDU) which consists of a multi-function keypad and LCD display. Pilots can call one of many display pages on the LCD, and pilots can compose and enter a variety of instructions to automated flight systems using the keyboard. While each crew member has their own CDU, for the most part the system does not support a shared representation, joint reference, and directed attention. When one crew member is programming the flight computers by working through multi-function controls and displays, as in the glass cockpits of newer commercial aircraft, it may be more difficult for the other crew member to pick up the other’s focus and actions without explicit attention or communication. One cannot see what the other is directing their attention to (except very crudely); one cannot physically direct the gaze of the other to an interesting part (at least not without a lot of spoken words); (3) one cannot see or assess the other agent’s activities preattentively. Note the contrast with the non-computerized representations analyzed by Hutchins (1991). Joint reference also can refer to human-machine coordination. For example, the pilot’s ability to see the ‘activities’ of automated systems is easily impaired by removing the external signs of those activities (Woods et al. 1994).

Auditory media such as voice loops in space control centers can also support the creation of a shared representation distributed across multiple cooperating agents. While outsiders tend to see the voice loops in space control centers as very noisy, especially when things get busy, space flight controllers find the voice loops quite important (Murray and Cox 1989). The voice loops possess the characteristics needed to support preattentive reference -- one can notice potentially interesting activities or events going on within the scope of responsibility of other controllers without drawing on limited attentional resources. This allows coordination of information and activity across multiple controllers with interacting areas of responsibility within the spacecraft and mission. In this case, preattentive reference directly supports the processes involved in joint reference.

Joint reference and a common frame of reference are part of specifying and designing coordinative structures: the architecture of roles across agents and the supporting tools for carrying out those roles coordinated with other agents, both machine and human.

5. Conclusion

The attention directing role of alarms is recognized superficially in the design of current alarm systems. Designers tend to think that by putting an ‘alarm’ on some condition they have ensured that the alarmed condition will function to break the operator away from other ongoing activities and switch their focus onto the condition alarmed. Similarly, if an individual signal, that in hindsight turns out to be important in an incident, is missed or misinterpreted by the practitioners, then regulators and engineering designers tend to think that they can mandate practitioner attention to that condition whenever it arises by turning up the perceptual salience of its associated alerting signal in isolation of other factors.

These approaches miss the fact that attentional processes function within a larger context that includes the state of the process, the state of the problem solving process, practitioner expectations, the dynamics of disturbance propagation. Considering each potentially anomalous condition in isolation and outside of the context of the demands on the practitioner will lead to the development of alarm and diagnostic systems that only exacerbate the alarm problem. For example, if changes tend to make the alarmed condition a ‘normal’ occurrence, then the alarm will become part of the routine flow of action, and it will not function to break attention away from other mental and physical activity. In aggregate, trying to make all alarms unavoidable redirectors of attention overwhelms the cognitive processes involved in control of attention and exacerbates the alarm problem. One kind of operational response to this should not really be surprising -- practitioners ignore or turn off the alarms (Sorkin and Woods 1985 showed that ignoring alarms is related to properties of the joint cognitive system; there is also a large body of field experience and field study results that show this response is symptomatic of poor support for directed attention, e.g., Cook et al. 1991; Moll van Charante et al. 1993). Even if the practitioner tries to switch attention every time an interrupt occurs, one must remember that there are costs associated with over-switching attentional focus -- loss of coherent situation assessment, failure to pick up suspended tasks or lines of reasoning, cognitive vagabonding (responding to every interrupt in isolation and never developing an integrated response strategy). Alarms are examples of attention directing cognitive tools. But one must recognize that directed attention is only meaningful with respect to the larger context of other activities and other signals.

Understanding action in the face of diverse, changing and highly uncertain situations depends critically on understanding how attentional processes are shaped by the tools available in the representation of the monitored process and on understanding how attentional control strategies shape the use or meaning of the interface mechanisms.

Acknowledgments

Research support was provided by the Aerospace Human Factors Research Division of the NASA Ames Research Center under Grant NCA2-351, Dr. Everett Palmer technical monitor. Additional support was provided by NASA Johnson Space Center under Grant NAG9-390, Dr. Jane Malin technical monitor.

References

Bower G. and Morrow, D. G. 1990, Mental models in narrative comprehension. *Science*, **24**, 44-48.

Broadbent, D. E. 1977, The hidden preattentive processes. *American Psychologist*, **32**, 109-118.

Bruner, J. 1986, *Actual Minds, Possible Worlds* (Harvard University Press, Cambridge MA).

Chase, W. and Simon, H. 1973, Perception in chess. *Cognitive Psychology*, **4**, 55-81.

Cook, R. I., Potter, S., Woods, D. D., and McDonald, J. S. 1991, Evaluating the human engineering of microprocessor controlled operating room devices. *Journal of Clinical Monitoring*, **7**, 217-226.

Cooper, G. E. 1977, *A Survey of the Status and Philosophies Relating to Cockpit Warning Systems* (NASA Ames Research Center, NASA-CR-152071, Moffett Field CA).

DeGroot, A. 1965, *Thought and Choice in Chess* (Mouton, The Hague).

Doyle, R. J., S. Sellers, S. and Atkinson, D. 1989, A focused, context sensitive approach to monitoring, in *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (Morgan-Kaufman, Los Angeles CA).

Doyle, R. J., Charest, L. K., Falcone, L. P. and Kandt, K. 1990, Addressing information overload in the monitoring of complex physical systems, in *Proceedings of the Fourth International Qualitative Physics Workshop* (ACM, NY).

Doyle, R. J., Chien, S., Fayyad, U. M. and Porta, H. 1992, Attention focusing and anomaly detection in real-time systems monitoring, in *Sixth Annual Workshop on Space Operations, Applications and Research* (National Aeronautical and Space Administration, Washington D.C.)

Federal Communications Commission, 1991, *Report on the September 17 1991 AT&T NY Power Outage at Thomas Street Switching Station and Network Disruption* (Common Carrier Bureau, Washington DC 20554).

Fink, R. T. 1984, *A Procedure for Reviewing and Improving Power Plant Alarm Systems* (Electric Power Research Institute, NP-3448, Palo Alto, CA).

Folk, C. L., Remington, R. W. and Johnston, J. C. 1992, Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology, Human Perception and Performance*, **18**, 1030-1044.

Gaver, W. 1991, Sound support for collaboration, in L. Bannon, M. Robinson and K. Schmidt (eds.)., *Proceedings of the Second European Conference on Computer-Supported Cooperative Work* (Amsterdam, The Netherlands, September 25-27).

Gopher, D. 1991,The skill of attention control, Acquisition and execution of attention strategies, in *Attention and Performance XIV* (Lawrence Erlbaum Associates, Hillsdale, NJ).

Hochberg, J.1986, Representation of Motion and Space in Video and Cinematic Displays, in K. R. Boff, L. Kaufman & J. P. Thomas (eds.)., *Handbook of Human Perception and Performance, Volume I* (John Wiley and Sons, New York).

Hutchins, E. 1991, *How a Cockpit Remembers Its Speed*. (Distributed Cognition Laboratory Report, Department of Cognitive Science, University of California, San Diego, August).

Hutchins, E. 1994,*Cognition in the Wild* (MIT Press, Cambridge, MA).

Kahneman, D. 1973, *Attention and Effort* (Prentice-Hall, Englewood Cliffs NJ).

Kemeny, J. G., et al. 1979, *Report of the President's Commission on the Accident at Three Mile Island* (Pergamon Press, New York).

Kragt, H. and 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.

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

Kubovy, M. and Pomerantz, J. R. (eds.). 1981, *Perceptual Organization* (Lawrence Erlbaum Associates, Hillsdale NJ).

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

Leroux, M. in press, Temporal reasoning in highly dynamic situations, in J.-M. Hoc, P. C. Cacciabue and E. Hollnagel (eds.)., *Expertise and Technology, Cognition and Human-Computer Cooperation* (Lawrence Erlbaum Associates, Hillsdale NJ).

Malin, J., Schreckenghost, D., Woods, D., Potter, S., Johannesen, L., Holloway, M. and Forbus, K. 1991, *Making Intelligent Systems Team Players: Case Studies and Design Issues* (NASA Johnson Space Center, TM-104738, Houston TX).

Mitchell, C. and Saisi, D. 1987, Use of model-based qualitative icons and adaptive windows in workstations for supervisory control systems. *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-17**, 573-593.

Moll van Charante, E., Cook, R. I., Woods, D. D., Yue, L. and Howie, M. B. 1993, Human-computer interaction in context, Physician interaction with automated intravenous controllers in the heart room, in H. G. Stassen, (ed.)., *Analysis, Design and Evaluation of Man-Machine Systems 1992* (Pergamon Press, New York).

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

Neisser, U. 1976, *Cognition and Reality* (W. H. Freeman, San Francisco).

Patterson, R. D. 1990, Auditory warning sounds in the work environment. *Philosophical Transactions of the Royal Society of London*, **B 327**, 485-492.

Pavard, B., Salembier, P., de Medeiros, E. and Benchekroun, T. 1989, *Simulation cognitive de la conduite de proces, Evaluation de differentes approaches* (Laboratoire d'Ergonomie, Conservatoire National des Arts et Metiers, ARD-AID T5403T, Paris France).

Potter, S. and Woods, D. D. 1991, Event-driven timeline displays, Beyond message lists in human-intelligent system interaction, in *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics* (IEEE, New York).

Rabbitt, P. 1984, The Control of Attention in Visual Search, in R. Parasuraman and D. R. Davies (eds.)., *Varieties of Attention* (Academic Press, New York).

Reiersen, C. S., Marshall, E. and Baker, S. M. 1988, An experimental evaluation of an advanced alarm system for nuclear power plants, in J. Patrick and K. Duncan (eds.)., *Training, Human Decision Making and Control* (North-Holland, New York).

Rochlin, G. I., La Porte, T. R. and Roberts, K. H. 1987, The self-designing high-reliability organization, Aircraft carrier flight operations at sea. *Naval War College Review*,, **Autumn**, 76-90.

Roth, E. M., Bennett, K. and Woods, D. D. 1987, Human interaction with an `intelligent' machine. *International Journal of Man-Machine Studies*, **27**, 479-525.

Sorkin, R. D. and Woods, D. D. 1985, Systems with human monitors, A signal detection analysis.*Human-Computer Interaction*, **1**, 49-75.

Sorkin, R. D., Kantowitz, B. and Kantowitz, S. 1988, Likelihood ratio alarms. *Human Factors*, **30**, 445-459.

Wolfe, J. M. 1992, The parallel guidance of visual attention. *Current Directions in Psychological Science*, **1**, 124-128.

Woods, D. D. 1994, Cognitive Demands and Activities in Dynamic Fault Management, Abduction and Disturbance Management, in N. Stanton (ed.)., *Human Factors of Alarm Design* (Taylor & Francis, New York).

Woods, D. D. 1995, Towards a Theoretical Base for Representation Design in the Computer Medium, Ecological Perception and Aiding Human Cognition, in J. Flach, P. Hancock, J. Caird, and K. Vicente (eds.)., *An Ecological Approach To Human Machine Systems I, A Global Perspective* (Lawrence Erlbaum, Hillsdale NJ).

Woods, D. D., Elm, W. C. and 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* (American Nuclear Society, New York).

Woods, D. D., Johannesen, L., Cook, R. I. and Sarter, N. 1994, *Behind Human Error, Cognitive Systems, Computers and Hindsight* (Crew Systems Ergonomic Information and Analysis Center, Dayton OH).

Woods, D. D., O'Brien, J. and Hanes, L. F. 1987, Human factors challenges in process control, The case of nuclear power plants, in G. Salvendy, (ed.)., *Handbook of Human Factors/Ergonomics* (Wiley, New York).

Figure 1. Illustration of a prototype display from an intelligent system for monitoring satellite communications. In the first panel, the system is quiet when it does not recognize any anomalies. In the second panel it has detected a loss of signal from the satellite, proposes possible hypotheses to account for this, and recommends actions for the human space controller to initiate. The label “red” indicates aspects of the display that change hue when the machine monitor detects a problem. (Peter Hughes, of NASA’s Goddard SpaceFlight Center graciously shared experience with this particular alarm problem.)

Figure 2. Illustration of a particular situation where the loss of signal from the satellite is expected -- a transponder switchover.