6

Alarm initiated activities

Neville Stanton

Introduction

The need to examine alarm handling behaviour stems from difficulties experienced by operators with industrial alarm systems (Pal and Purkayastha, 1985). Across a range of industrial domains, alarm systems appear to place the emphasis on detection of a single event, rather than on considering the implications of the alarm within the task (Stanton, 1993). Therefore, current industrial systems do not appear to make optimum use of human capabilities which could improve the overall human supervisory control performance (Sorkin, 1989). This is desirable because we are unlikely to remove human operators from the system. This would require a level of sophistication not possible in the foreseeable future. However, the reluctance to leave a machine in sole charge of 'critical' tasks is likely to mean that human operators will still be employed in a supervisory capacity because of concern about break-down, poor maintenance, as well as ethical concerns. Therefore we need to capitalize on the qualities that operators bring to the 'co-operative endeavour' of human-machine communication. Alarm problems are further confused by the inadequacies of peoples' understanding of what constitutes an 'alarm' (Stanton and Booth, 1990). Most definitions concentrate on a subset of the qualities or properties, for example 'an alarm is a significant attractor of attention' or 'an alarm is a piece of information'. In fact, an alarm may be considered from various perspectives (Singleton, 1989), which need to be integrated into one comprehensive definition if the term is to be understood in its entirety. An 'alarm' should be defined within a systems model and consider how each of the different perspectives contribute to the interpretation of the whole system (Stanton, Booth *et al.,* 1992). In this way, one may examine the role of the human operator in response to alarm information, in order to develop a model of alarm handling that will ultimately influence alarm system design. A model may be considered to be a description or representation of a process that enables analysis of its form to be undertaken. A model of alarm handling is necessary to guide research, so that we may ask appropriate questions and utilize suitable empirical techniques to yield answers.

The development of models to understand human behaviour within complex systems is not a new endeavour (Edwards and Lees, 1974; Broadbent, 1990). It has been the domain of cognitive psychologists and human factors researchers alike. Models serve practical purposes, such as:

- a framework to organize empirical data;
- a prompt for investigation;
- to aid design solutions;
- to compare with actual behaviour;
- to test hypotheses and extrapolate from observable inferences;
- to measure performance;
- to force consideration of obscure or neglected topics.
  (Pew and Baron, 1982).

Models may be coarsely split into two types: quantitative and qualitative. Quantitative models are computational, (for example: simulations and ana-lytic or process models) whereas qualitative models are descriptive. Quantitative models can produce mathematically precise estimates of performance (Broadbent, 1990; Elkind, Card *et al.,* 1990), but they are limited to use in highly specialized and restricted domains. Often the lack of hard data to put into a quantitative model of human behaviour means that one must first develop qualitative models. These serve as a basis for collecting the necessary empirical data that could eventually provide the information for a quantitative model.

Many qualitative models of human intervention in control room incidents have been proposed (Edwards and Lees, 1974; Rasmussen, 1976; Rouse, 1983; Hale and Glendon, 1987; Swain and Weston, 1988). The best known of these are the models of Rouse (1983) and Rasmussen (1976, 1983, 1984, 1986). Rasmussen's Skill-Rule-Knowledge (SRK) framework is extensively cited in the literature, and has been accepted as 'the industry standard' (Reason, 1990). The SRK framework distinguishes between three levels of performance that correspond with task familiarity. At the lowest level, skill-based performance is governed by stored patterns of proceduralized instructions. At the next level, behaviour is governed by stored rules, and at the highest level, behaviour is governed by conscious analytical processes and stored knowledge. Pew, Miller *et al.* (1982) comment on the strengths of Rasmussen's framework which they present as a decision making model which contains three essential elements that are consistent with human problem solving: data processing activities, resulting states of knowledge and shortcuts in the 'stepladder' model (discussed next).

Reason (1990) commented on Rasmussen's eight stages of decision making for problem solving: activation, observation, identification, interpretation, evaluation, goal selection, procedure selection and activation. He suggested that Rasmussen's major contribution was to have charted the shortcuts that human decision makers take in real situations (i.e. the stepladder model) which result in 'highly efficient, but situation-specific stereotypical reactions'. Pew and Baron (1982) provides an example of problem detection, for which the operator collects limited data and may immediately conclude that a specific control action must be executed (skill-based behaviour). Alternatively, the operator may additionally identify the system state and then select and execute a procedure that results in an action sequence (rule-based behaviour). Finally when the circumstances are new or the specific combination of circumstances does not match known ones, then the whole range of problem solving behaviour is called forth (knowledge-based behaviour). Reason (1988b) suggests that most incidents are likely to require this last type of behaviour, because although they may start in a familiar way they rarely develop along predictable lines. It is this unpredictable development that gives the greatest cause for concern, particularly when the true nature of the incident departs from the operator's understanding of it (Woods, 1988). As Reason (1988b) notes:

> each incident is a truly novel event in which past experience counts for little, and where the plant is returned to a safe state by a mixture of good luck and laborious, resource limited, knowledge-based processing.

From an extensive review of the literature on failure detection, fault diagnosis and correction, Rouse (1983) identified three general levels of human problem solving, namely:

- recognition and classification;
- planning; and
- evaluation and monitoring.

Within each of these levels Rouse assigns a three stage decision element to indicate whether the output of each stage is skill-based, rule-based or knowledge-based, rather like Rasmussen's framework. Firstly it is assumed that the individual is able to identify the context of the problem (recognition and classification), and then is able to match this to an available 'frame'. If a 'frame' does not exist then the individual has to resort to first principles. At the planning level, the individual must decide if a known procedure can be used, or whether alternatives have to be generated. Problem solving is generated at the lowest level where plans are executed and monitored for success. Familiar situations allow 'symptomatic' rules (i.e. rules based upon identifying familiar plant symptoms), whereas unfamiliar situations may require 'topographic' rules (i.e. rules based upon an understanding of the physical topography of the plant and the cause-effect relationships of the components). However, it has been argued that human problem solving is characterized by its opportunistic nature, rather than following a hierarchical information flow (Rouse, 1983; Hoc, 1988), with all levels being employed simultaneously. This would suggest a problem-solving heterarchy utilizing parallel processing. Therefore, the SRK model is not without its critics. Bainbridge (1984) suggests that at best it presents an oversimplified account of cognitive activity, and that at worst the inferences drawn may be wrong. Her main criticisms may be summarized as:

- a confusion of the terminology;
- a failure to represent all aspects of human behaviour;
- missing important aspects for the understanding of human cognition.

She warns of the danger of a strict application of the SRK framework which might restrict the flexibility of human behaviour, for example, by providing displays that can only be used for limited purposes. However, she does accept that it provides the basic idea of cognitive processes. Most of the criticism of the SRK framework has arisen either from a misunderstanding of the original intention, which was to provide a framework rather than a grand psychological theory, or from inappropriate application (Goodstein, Andersen *et al.,* 1988). Thus within its accepted limitations, it has remained robust enough to be considered a working approximation to human cognitive activities and allows for some prediction and classification of data.

Much of the attention paid to the SRK framework has been in the domain of human supervisory control, and Reason (1988b) presented the 'catch-22' of such systems.

- The operator is often ill-prepared to cope with emergencies, because the relatively low frequency of the event means that it is likely to be outside his/her experience. Moreover, high levels of stress are likely to accompany the emergency, making the operator's task more difficult.
- It is in the nature of complex, tightly-coupled, highly interactive and partially understood process systems to spring nasty surprises (Perrow, 1984).

The first point was made eloquently by Bainbridge (1983) in her discussion of the 'ironies of automation'. In the design of complex systems, engineers leave the tasks they cannot automate (or dare not automate) to the human, who is meant to monitor the automatic systems, and to step in and cope when the automatic systems fail or cannot cope. However, an increasing body of human factors knowledge and research suggests that the human is poor at monitoring tasks (Moray, 1980; Wickens, 1984; Moray and Rotenberg, 1989). When the humans are called to intervene they are unlikely to do it well. In other words, removing the humans from control is likely to make the task harder when they are brought back in (Hockey, Briner *et al.,* 1989). It has been suggested that diagnosis and control behaviour and quite different (Wickens, 1984). However, diagnosis behaviour is likely to be (at least in part) adapted to the way in which the information is presented to the operator and vice versa. Therefore emphasis needs to be put on understanding how

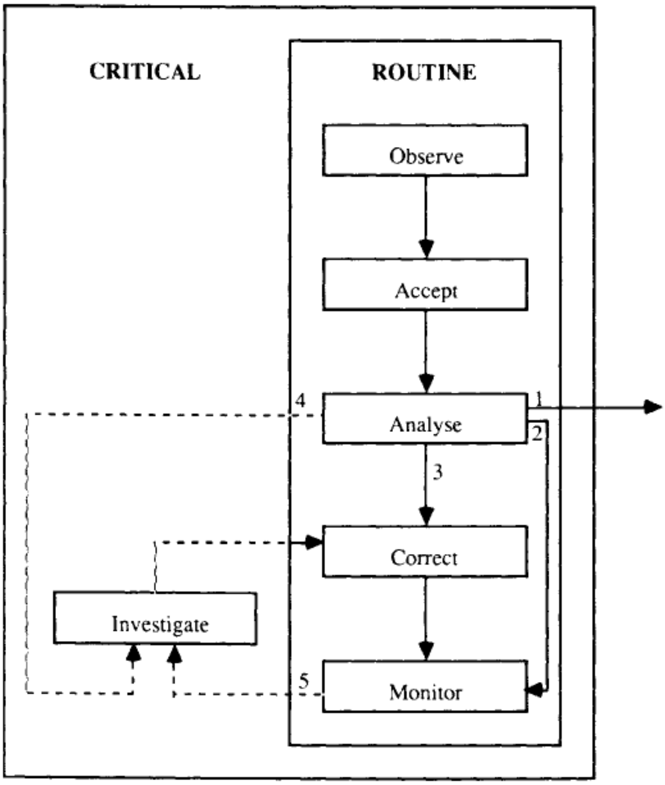

Figure 6.1 Model of alarm initiated activities.

the operator uses and processes the information, and to relate this understanding back to human cognitive activity in fault management in general.

Model of alarm initiated activities

The following model was constructed by Stanton (1992). As shown in Figure 6.1, it highlights the difference between routine incidents involving alarms (plain lines) and critical incidents involving alarms (dotted lines). The distinction between 'routine' and 'critical' is determined by the operator in the course of alarm handling. Although there are common activities to both types of incident (Figure 6.1), critical incidents require more detailed investigations. It is proposed that the notion of alarm initiated activities (AIA) is used to describe the collective of these stages of alarm event handling. The term 'activities' is used here to refer to the ensuing cognitive modes as well as their corresponding behaviours, both of which are triggered by alarms. The AIA are assumed to be distinctly separate activities to 'normal' operation in supervisory control tasks.

Table 6.1 Example of alarm initiated activities

Event	Outcome	AIA
1 Pump temperature exceeds alarm threshold.	'Pump ABC Temp High' alarm flashes accompanied by tone	
2 Operator hears alarm tone	Operator looks at alarm panel	Observe
3 Operator presses 'alarm accept' key	Alarm stops flashing and tone is silenced	Accept
4 Operator reads alarm message	(a) ignore alarm (b) monitor situation (c) reset alarm (d) investigate	Analyse
5 Operator investigates cause of pump ABC overheating	Operator finds valve XYZ closed	Investigate
6 (a) Operator opens valve XYZ (b) Operator stops pump ABC	(a) Valve XYZ opens (b) Pump ABC stops	Correct
7 Operator intermittently checks pump ABC	Pump ABC temperature eventually comes below threshold for 'Temp High' alarm	Monitor
8 Operator resets 'Pump ABC Temp High' alarm	'Pump ABC Temp High' alarm returns to non-active state	Correct

Typically control desk engineers (CDEs) report that they will observe the onset of an alarm, accept it and make a fairly rapid analysis of whether it should be ignored (route 1), monitored (route 2), dealt with superficially (route 3) or require further investigation (route 4). Then, even if they feel that it may require further investigation, they may still try to correct and cancel it (route 3) just to see what happens. If it cannot be cleared, then they will go into an ivestigative mode to seek the cause (route 5). Then in the final stage the CDEs will monitor the status of the plant brought about by their corrective actions. The need to resort to the high cognitive level 'investigation' is what distinguishes critical from routine incidents. The stages of activity may be considered with the help of an example of alarm handling taken from a manufacturing industry (Table 6.1).

Consider the filling of a tank from a storage vessel through a pipe with a valve and pump in-line. The operator in the control room is busy with various aspects of the task, such as the setting up of equipment further on in the process when he/she hears an audible alarm (event 2 in Table 6.1). The alarm is acknowledged by the cancellation. The operator now has a variety of options, as it is not yet known why the alarm telling the operator that the pump has overheated was triggered. There are a number of plausible explanations, such as:

1. there is a physical fault with the pump;
2. the storage vessel is empty;

3. the supply pipe is blocked or leaking; or
4. the valve is closed.

Given these degrees of uncertainty, there are several different remedial actions open to the operator as shown by outcomes to event 4. One path to saving the pump might be to stop it running (event 6b). Alternatively the operator may attempt to find the cause of overheating, which may be due to the valve not being opened before the pump was switched on. This may lead the operator to open the valve (event 6a) and then intermittently check the status of 'pump ABC' (event 7). Eventually the alarm will change status and enable the operator to reset it (event 8).

The above is an idealized description of a successful path through the series of events, and as such gives a simplified account of the true nature of the task. It assumes that the operator was successfully able to identify the reason for the alarm, although the alarm cue did not directly point to it. In this case there was a variety of plausible alternatives, each of which would require investigation. Whether or not exhaustive discounting actually takes place depends on the operator being able to bring them to mind.

The criteria for defining success are also ambiguous. If the operator stops the pump (event 6b), this would lead to the alarm being cleared, thus providing the opportunity to route the product through another pipe to fill the tank. Such a strategy would, perhaps, have been equally successful as the first alternative selected. In reality there may be many different possible courses of action competing for the operator's time and attention depending on the number of active alarms. The task is made even more difficult by the fact that alarms may also be grouped by events, and be interdependent on each other. This is particularly true in closely coupled systems (Perrow, 1984) with feedback loops. Such grouping can make the task of distinguishing cause and effect very difficult and, in turn, add to the inherent ambiguities described earlier.

As the example demonstrates, an alarm handling sequence can be described as consisting of a number of generic activity stages. The activities are illustrated in the AIA (alarm initiated activities) column of Table 6.1. Studying the alarm handling activities employed by operators might give some indication of how best to design alarm systems. This argument will be developed within the chapter.

Therefore, a consideration of the literature is required to make further inference about the requirements of these stages of handling. These AIAs will provide the framework of the review and guide subsequent research. The review is presented in the following sections: observe, accept, analyse, investigate, correct and monitor.

Observe

The observe mode is characterized by the initial detection of abnormal plant conditions. Detection is the act of discovering any kind of undesired deviation(s) from normal system operations (Johannsen, 1988). Bainbridge (1984) suggests that there are three main ways of detecting abnormal plant conditions:

- responding to an alarm;
- thinking of something that needs to be checked;
- incidentally noticing that something is wrong whilst attending to somethingelse.

Failure to detect an abnormal situation may occur for a number of reasons (Moray, 1980):

- the relevant variable is not displayed;
- the signal to noise ratio is too low;
- the expectation of the operators leads to a misinterpretation of the information;
- the information may be ignored due to attention being directed on other variables;
- there may be too much information.

Under normal conditions Moray suggests that most systems are adequate to allow visual scanning to support monitoring tasks. However, when very rapid changes occur the task becomes very difficult. Prolonged activity of this kind is likely to reduce the efficiency of human cognitive activities as

> several concurrent activities may compete for access to a particular (cognitive) 'resource'…the cost of errors may be very great.
>
> Hockey, Briner *et al.* (1989)

Counter to an intuitive notion of the control task, Moray (1980) suggests that the better the system is known to an operator, the less likely he/she will discover an abnormal state. He implies that this is due to the reliance of the operator on past experience and the correlation between variables to predict future states. This leads to a failure to observe current values. Therefore abnormal values are undetected. This proposition is similar to the observations of Crossman and Cooke (1974) who noticed that skilled tracking behaviour was primarily 'open-loop'. Tracking is compensatory (that is it occurs after the event), therefore when dealing with highly familiar data the human is likely to fill in the gaps or miss the data. Reason (1990) suggests that as fault detection moves from being knowledge-based to becoming skill-based, it is likely to suffer from different types of error. Reason proposes that skill-based behaviour is susceptible to slips and lapses whereas knowledge-based behaviour is susceptible to mistakes.

In a series of experiments aimed at investigating fault detection in manual and automatic control systems, Wickens and Kessel (1981) concluded that automating the system does not necessarily reduce the mental workload of the human controller. Firstly they noticed a paradox of task operation. In manual control, operators are able to continually update their 'model' of the system, but are also required to perform two tasks: control and detection. Whereas in automatic control they had only the detection task, but were not 'in-loop' to update their 'model'. This means that removing the human from the control loop may reduce the attention paid to the system state. Wickens and Kessel suggest that whether the manual or automatic control task performance was superior would depend largely upon the relative workload, i.e. under some conditions workload might favour manual control and in others workload might favour automatic control. Automation shifts the locus of the information processing demands. In manual control, the emphasis is primarily on 'responding', whereas in automatic control the demands are primarily located in 'perception' and 'central processing'. Under the SRK framework the shift is from skill-based behaviour to knowledge- and rule-based behaviour.

Wickens and Kessel also suggest a 'fragility' of failure detection performance as:

- it cannot benefit from borrowed resources of responding;
- it deteriorates when responding demand is increased.

In summary, it appears that detection has the 'worst of both worlds'. This may represent an intrinsic characteristic of detection tasks in general.

In a series of investigations into fault management in process control environments, Moray and Rotenberg (1989) observed that subjects:

- display cognitive lockup when dealing with a fault;
- prefer serial fault management;
- experience a time delay between noticing a fault and dealing with it.

Moray and Rotenberg noticed that when dealing with one fault their subjects would not take action on another. This is linked to the preference for dealing with faults serially, rather than concurrently. Moray and Rotenberg were however, unable to distinguish between cause and effect, i.e. whether cognitive lockup leads to subjects dealing with faults serially or vice versa. In process systems, serial fault management may not produce optimum process performance, but it may make task success more likely, as interruptions in fault management (to deal with other faults) may cause the human operator to forget important aspects of the first task that was being worked on. The data collected by Moray and Rotenberg can explain the time delay between looking at a fault and dealing with it. The data showed that a fault is examined many times before intervention is initiated. Their eye-movement data demonstrate that just because operators are not actively manipulating controls we cannot assume that their task load is low. Moray and Rotenberg's data suggest that the operator is actively processing information even in apparently non-active periods. They claim that an operator might observe an abnormal value, but fail to take action for at least three reasons:

- the evidence was not strong enough to lead to a diagnosis for appropriate action;

- the operator was already busy dealing with another fault and wishes to finish that problem before starting a new one;
- although the abnormal value was observed, it was not perceived as abnormal.

They conclude from their data that the second of these proposals appears most likely in their investigation. The locking-up of attention is a phenomenon that has been repeatedly reported in the literature (e.g. Moray and Rotenberg, 1989; Hockey, Briner *et al.,* 1989; Wickens, 1984) and appears to be a intrinsic characteristic of human cognitive processing. As Wickens (1984) expresses it:

> …it is reasonable to approximate the human operator as a single-channel processor, who is capable of dealing with only one source of information at a time.

The irony of attracting the operator's attention to the new alarm information is that successful attraction will necessarily mean distracting the operator from other aspects of the task. The interruption may not be welcome as it may interfere with some important operation. Therefore the alarm system needs to show that a problem is waiting to be dealt with, rather than forcing the operator to deal with it unless the alarm merits immediate action, and enable the operator to distinguish between alarms that relate to separate events. Moray and Rotenberg (1989) report that the probability of looking at a fault and dealing with it may be described in terms of a logarithmic relationship between probability of detection and time since its occurrence.

Accept

The acceptance of an alarm is taken to be acknowledgement or receipt. This is normally a physical action that takes the alarm from its active state to a standing state. Jenkinson (1985) proposed that audible and visual cues should be combined to reduce the visual search task, as the operator has to move within the workspace, and visual information alone is insufficient. Normally the receipt of an alarm is accompanied by the silencing of the audible cue, and a change in some aspect of the visual coding, such as from flashing to illuminated. However, this change in visual and auditory state may make it difficult to tell when an alarm has been accepted. For example, in an annunciator or mimic display, once the flashing code has stopped there may be no means of recording the time or order of occurrence of the alarm. So by accepting it, the operator loses some information about the alarm that may be essential for the subsequent AIAs, (such as 'analyse' or 'investigate') to be performed effectively. However, the alarm may be considered to be in one of four possible states:

- not activated;
- activated but not accepted;

- accepted but not reset;
- reset.

Resetting an alarm is the acknowledgement by the operator that the initiating condition is no longer present. It extinguishes the alarm, returning it to its first state: not activated. The indication that an alarm is waiting to be reset is normally in the form of a marker or code (Jenkinson, 1985) to inform the operator of its new state.

The designers of alarm systems have to consider whether to allow group acknowledgement of alarms, or to insist on each alarm being acknowledged individually. Unfortunately the literature is inconclusive. Group acknowledgement of alarms may cause the operators to deal inadvertently with a signal (Kragt and Bonten, 1983) but single acknowledgement may fare no better (Kortlandt and Kragt, 1980). With group acknowledgement it is possible that the operator could miss a signal by accepting *en masse* and scan the alarm list or matrix. However, in periods of high alarm activity it is likely that single acknowledgement actions will resemble group acknowledgement, as the operator repeatedly presses the 'accept' key without reading the alarm message (Stanton, 1992). Reed and Kirwan (1991), however, describe the development of an alarm system that requires operators to accept each alarm individually.

Under certain operational situations up to 200 alarms could be presented. They claim that the simplicity of the task will mean that single acknowledgement of each of the 200 alarms will not be unduly problematic. What they do not acknowledge is that tying the operators up in this simple acceptance task prevents them from moving further on in the alarm initiated activities. This could become a problem if there are other critical failures within the process that are hidden within the 200 alarms presented. Further, an operator may sometimes accept a signal just to get rid of the audible signal (Kragt and Bonten, 1983; Sorkin, 1989). This presents a paradox in design, because the operator is made aware of a change in the process state by the presence of the signal attracting attention. Failure to attend to the alarm will mean that it is impossible to pass this information on to the subsequent stages of AIAs. Masking of a fault may result from too many alarms. This was the most often cited reason for missing alarms in recent studies (Stanton, 1993).

Analyse

Analysis may be considered to be the assessment of the alarm within the context of the task that is to be performed and the dynamics of the system. Analysis appears to involve a choice of four options (ignore alarm, monitor situation, deal with alarm superficially or investigate cause) and therefore involves some rudimentary search of context to reach an appropriate judgement. Easterby (1984) proposed that a variety of psychological processes are used by an operator in control of a machine, such as: detection, discrimination, identification, classification, recognition, scaling, ordering and sequencing. He suggested that the control panel may be considered as a map of the operator's task:

> the display must therefore define the relationships that exist between the machine elements, and give some clues as to what to do next.

This is essentially the operator's task in analysis: to decide what to do next. Operators are often required to search for the relevant information to base their decisions on, as in VDU-based control systems the information is not necessarily available immediately, and can only be obtained after request (Kragt and Bonten, 1983).

From the reported behaviours of plant operators, the results of the analysis stage of AIAs determine the future course of action: ignoring the alarm, monitoring the system, making superficial corrective actions to cancel the alarm, or going into an investigative mode. This puts an emphasis on the alarm to convey enough information to make this decision without involving the operators in too much effort as there may be other demands upon their attention. To some extent operators may be aided in the task by a current awareness of the plant state. For example, if they know that a part of the plant is in maintenance, then they are unlikely to be surprised that the value of a particular variable is outside its normal threshold. Alternatively if they are tracking the development of an incident, an alarm may confirm their expectations and therefore aid diagnosis. However, it is also possible that the operators may wrongly infer the true nature of the alarm leading to an inappropriate analysis and subsequent activity. It is important to note that the presence of the alarm by itself may not directly suggest what course of action is required, but only reports that a particular threshold has been crossed. In the search for the meaning of the alarm, the manner in which it is displayed may aid or hinder the operator. For example alarm lists show the order in which the alarm occurred; alarms within mimic displays map onto the spatial representation of the plant, and annunciator alarms provide the possibility for pattern recognition.

These different ways of presenting alarm information may aid certain aspects of the operator's task in analysis, such as indicating where the variable causing the alarm is in the plant; what the implications of the alarm are; how urgent the alarm is, and what should be done next. Obviously different types of information are conveyed by the different ways to present alarm information mentioned (lists, mimics and annunciators). The early classification process may be enhanced through pairing the visual information with auditory information such as tones or speech. Tones are abstract and would therefore require learning, but may aid a simple classification task such as urgency (Edworthy and Loxley, 1990).

Tones provide constant information and are therefore not reliant on memory for remembering the content of the message. They are reliant on memory for recalling the meaning of the message. Whereas speech is less abstract and rich in information, it is varied and transitory in nature, so whilst it does have the possibility of providing complex information to the operator in a 'hands-free eyes-free' manner, it is unlikely to find favour as an alarm medium in process control (Baber, 1991).

It has been speculated that text and pictures are processed in a different manner (Wickens, 1984), and there are alternative hypotheses about the underlying cognitive architectures (Farah, 1989). Wickens' dual face multiple resource theory and stimulus-cognitive processing-response (SCR) compatibility theory offer an inviting, if mutually irrefutable, explanation of information processing. Wickens' theories predict that the modality of the alarm should be compatible with the response required provided that the attentional resources for that code are not exhausted. If attentional resources for that code are exhausted, then another input modality that does not draw on the same attentional resources should be used. Despite the attraction of Wickens' explanation, based on a wealth of data involving dual task studies, there is still some contention regarding the concept of separate information processing codes. Farah (1989) draws a clear distinction between the three main contending theoretical approaches to the representation of peripheral encoding and internal cognitive processing. First, Farah suggests that although encoding is specific to the input modality, internal processing shares a common code. Second, the single code approach is favoured by the artificial intelligence community, probably because of the computational difficulties of other approaches (Molitor, Ballstaedt *et al.,* 1989). Alternatively (third) the 'multiple resource' approach proposes separate encoding and internal processing codes (Wickens, 1984). Farah (1989) suggests that recent research points to a compromise between these two extremes.

Recent studies have shown that a combination of alphanumeric and graphic information leads to better performance than either presented alone (Coury and Pietras, 1989; Baber, Stammers *et al.,* 1990) It might similarly be speculated that the combination of codes in the correct manner may serve to support the analysis task. The model of AIAs implies that different aspects of the code might be needed at different points in the alarm handling activity. Thus the redundancy of information allows what is needed to be selected from the display at the appropriate point in the interaction. The type of information that is appropriate at any point in the interaction requires further research.

Investigate

The investigative stage of the model of AIAs is characterized by behaviour consistent with seeking to discover the underlying cause of the alarm(s) with the intention of dealing with the fault. There is a plethora of literature on fault diagnosis, which is probably in part due to the classical psychological research available on problem solving. The Gestalt psychology views provide an interesting but limited insight into problem solving behaviour, confounded by vague use of the terminology. Research in the 1960s was aimed at developing an information processing approach to psychology in general, and to problem solving in particular, to:

> …make explicit detailed mental operations and sequences of operations by which the subject solved problems.
>
> Eysenck (1984)

A closer look at research from the domain of problem solving illustrates this clearly. Problem solving may be considered analogous to going through a maze, from the initial state towards the goal state. Each junction has alternative paths, of which one is selected. Moving along a new path changes the present state. Selection of a path is equivalent to the application of a number of possible state transforming operations (called operators). Operators define the 'legal' moves in a problem solving exercise, and restrict 'illegal' moves or actions under specific conditions. Therefore a problem may be defined by many states and operators, and problem solving consists of moving efficiently from our initial state to the goal state by selecting the appropriate operators. When people change state they also change their knowledge of the problem. Newell and Simon (1972) proposed that problem solving behaviour can be viewed as the production of knowledge states by the application of mental operators, moving from an initial state to a goal state. They suggested that problem solvers probably hold knowledge states in working memory, and operators in long term memory. They problem solver then attempts to reduce the difference between the initial state and the goal state by selecting intermediary states (subgoals) and selecting appropriate operators to achieve these. Newell and Simon suggest that people move between the subgoal states by:

- noting the difference between present state and goal state;
- creating a subgoal to reduce the difference; and
- selecting an operator to achieve this subgoal.

Thus it would appear that the cognitive demand of the task is substantially reduced by breaking the problem down, moving towards the goal in a series of small steps. A variety of computer-based systems have been produced in an attempt to model human problem solving, but none have provided a wholly satisfactory understanding. This is not least because they are unable to represent problem solving in everyday life, and computer models rely on plans, whereas actions may be performed in a number of ways. As Hoc (1988) proposes:

> A problem will be defined as the representation of a task constructed by a cognitive system where this system does not have an executable procedure for goal attainment immediately at its disposal. The construction of a task, representation is termed understanding, and the construction of the procedure, problem solving.

This means that the same task could be termed a problem for some people, but not for others who have learned or developed suitable procedures (Moran, 1981). The difficulty in analysing problem solving is the human ability to perform cognitive activity at different levels of control at the same time. Rasmussen's SRK framework is useful in approximating these levels, but the entire activity leading to a goal can seldom be assigned to one, and usually occurs at all levels simultaneously. Hoc (1988) sees problem solving as involving two interrelated components: problem understanding (the construction of a coherent representation of the tasks to be done) and procedure searching (the implementation of a strategy to find or construct a procedure). This suggests that there is an 'executive controller' of the problem solving activities which directs the choices that are taken (Rouse, 1983). Planning is the guiding activity that defines the abstract spaces and is typically encountered in problem solving. Hoc (1988) believes that planning combines top-down components (creating new plans out of old ones) with bottom-up components (elaborating new plans or adapting old plans). Thus he suggests that an information representation that supports the shift between these components would result in more efficient strategies. Human factors is essentially about the design of environments that suit a wide range of individuals. Therefore presentation of information that only suits one strategy, or particular circumstances, is likely to frustrate the inherent variation and flexibility in human action.

Landeweerd (1979) contrasts diagnosis behaviour with control, proposing that, in control, the focus of attention is upon the forward flow of events, whereas diagnosis calls for a retrospective analysis of what caused what. Wickens (1984) widens the contrast by suggesting that the two tasks may be in competition with each other for attentional resources and that the two phases of activity may be truly independent. However, whilst diagnosis certainly does have a retrospective element in defining the problem, it certainly has a forward looking element of goal directed behaviour in correcting the fault. Landeweerd (1979) suggests that the type of internal representation held by the operator may predict control behaviour. Although his findings are tentative they do suggest that different types of information are used in problem search and problem diagnosis. During search only the mental image (i.e. a mental picture of the plant) plays a role, whereas the mental model (i.e. an understanding of the cause-effect relationships between plant components) plays a more important role in diagnosis. Landeweerd explains that this is because search behaviour is working from symptoms to causes, whilst diagnosis relates the results from the search activities to probable effects. However, the correlations between the mental image and mental model data obtained by Landeweerd were not very high, and the internal representations may be moderated by other variables, such as learning or cognitive style.

A number of studies have suggested that the type of knowledge acquired during problem solving may indicate success in dealing with failures. In a comparison of training principles with procedures, the results indicate that rule-based reasoning is better for routine failures, whereas knowledge-based reasoning is better for novel situations (Mann and Hammer, 1986; Morris and Rouse, 1985). Rouse and Rouse (1982) suggest that selection of strategies for problem solving tasks could be based upon cognitive style as certain styles may reflect more efficient behaviour. However, the results of further work indicate that the variations found in individuals highlight the need for more flexible training programmes.

In an analysis of the convergence or divergence of hypothesis testing in problem solving, Boreham (1985), suggests that success may be enhanced by the subject considering more hypotheses than absolutely required. This suggestion implies that a certain redundancy in options available may aid the task of problem solving by getting the subject to consider the problem further in order to justify their choice of intervention strategy. However, Su and Govindaraj (1986) suggest that the generation of a large set of plausible hypotheses actually degrades performance due to the inherent limitations of information processing ability. Providing many possible alternatives, therefore, makes the identification of the correct alternative more difficult, whereas a limited selection would presumably make the decision task easier.

Brehmer (1987) proposes that the increasing complexity of system dynamics makes the task of fault management more one of utilizing diagnostic judgment in a situation of uncertainty and less one of troubleshooting. The supervisory control task is becoming more like that of a clinician in diagnosing various states of uncertainty rather than the application of troubleshooting methods such as split-half strategies. Research on the diagnostic process suggests that the form of judgment tends to be simple (little information used, and it tends to be used in an additive rather than configurational way); the process is generally inconsistent, there are wide individual differences and individuals are not very good at describing how they arrived at judgments (Brehmer, 1987).

> The problem of fault diagnosis in complex systems arrives not from major catastrophic faults, but from cascades of minor faults that together overwhelm the operator, even though none would do so singly.
>
> Moray and Rotenburg (1989)

Thus the nature of the process plant may be considered to be greater than the sum of its parts due to the: inter-relation of the parts of the process plant, the system dynamics, many feedback loops and the inherent ambiguity of the information for diagnostic evaluation (Moray, 1980). This change in the nature of the task from troubleshooting to diagnostic judgement in a situation of uncertainty has implications for the way in which information is presented. As Goodstein (1985) suggests, this needs to change also. Goodstein proposes that the information should move away from the traditional physical representation of plant components toward a functional representation as, he suggests, this is closer to the operators' understanding of the plant. Thus the functional representation requires less internal manipulation.

Moray and Rotenberg's (1989) investigation into fault management in process control supported the notion that humans inherently prefer to deal with faults serially, rather than by switching between problems. They claim that this has serious implications for fault management in large complex systems, where any response to faults occurring late in the sequence of events would be greatly delayed, even if the later faults were of a higher priority than the earlier faults. It has been further proposed that in dealing with complex systems, humans are susceptible to certain ‘primary mistakes’. These include: an insufficient consideration of processes in time, difficulties in dealing with exponential events and thinking in terms of causal series rather than causal nets (Reason, 1988c). These factors combined may help explain why the operators’ understanding of the system state may not always coincide with the actual system state (Woods, 1988). Clearly the investigative task is very complex, and a means of representation to aid the operators’ activities needs to consider the points mentioned here.

Correct

Corrective actions are those actions that result from the previous cognitive modes in response to the alarm(s). In a field study, Kortland and Kragt (1980), found that the limited number of actions that followed an alarm signal suggested that the main functions of the annunciator system under examination were to be found in its usefulness for monitoring. This supports Moray and Rotenberg’s (1989) assertions that low observable physical activity is not necessarily accompanied by low mental activity. The majority of signals analysed by Kortland and Kragt (1980) were not actually ‘alarms’ in the sense that a dangerous situation was likely to occur if the operator did not intervene, and this must have led to its use as a monitoring tool, which has also been observed in other studies (Kragt and Bonten, 1983). However, they found that during periods of high activity the operator may pay less attention to individual signals, and mistaken actions could occur. Thus, lapses in attention in early AIA modes may lead to inappropriate corrective actions. The choice of compensatory actions is made by predicting the outcome of the alternatives available, but these evaluations are likely to be made under conditions of high uncertainty (Bainbridge, 1984). Bainbridge offers eight possible reasons for this uncertainty in the operator:

- action had unpredictable or risky effects;
- inadequate information about the current state of the system;
- wrong assumption that another operator had made the correct actions;
- precise timing and size of effects could not be predicted;
- no knowledge of conditions under which some actions should not be used;
- no knowledge of some cause-effect chains in the plant;
- difficulty in assessing the appropriateness of his/her actions;
- distractions or preoccupations;

It is assumed that knowledge embodied in the form of a coherent representation of the system and its dynamics (i.e. a conceptual model) would facilitate control actions, but the evidence is not unequivocal (Duff, 1989). Reason (1988a) suggests, in an analysis of the Chernobyl incident, that plant operators operate the plant by 'process feel' rather than a knowledge of reactor physics. He concludes that their limited understanding was a contributing factor in the disaster. However, under normal operation the plant had given service for over three decades without major incident. It was only when their actions entered into high degrees of uncertainty (as listed by Bainbridge, 1984) and combined with other 'system pathogens' that disaster became inevitable (Reason, 1988a).

Open-loop control strategies appear to be preferable in process control because of the typically long time constants between an action being taken and the effect of that manipulation showing on the display panel. Under such circumstances, closed-loop process manipulation might be an inefficient and potentially unstable strategy (Wickens, 1984). Under consideration of the 'multiple resources' representation of information processing, Wickens (1984) proposes that 'stimulus-cognitive processing-response' (SCR) compatibility will enhance performance, and conversely 'SCR' incompatibly would be detrimental to performance. This relationship means that the alarm display needs to be compatible with the response required of the operator. This framework may be used to propose the hypothetical relationship between alarm type and compatible response. This may be summarized as: text and speech based alarms would require a vocal response, whereas mimic and tone based alarms would require a manual response. Annunciator alarms appear to have both a spatial and a verbal element. Presumably they could, therefore, allow for either a verbal or a manual response. This last example highlights some difficulties with the SCR compatibility idea. Firstly, just because an input modality appears to be either verbal or spatial it does not necessarily allow for a simple classification into an information processing code. Secondly, many real life situations cross both classifications. Thirdly, control rooms usually require some form of manual input, and speech based control rooms, although becoming technically feasible, may be inappropriate for some situations (Baber, 1991a). Finally, Farah (1989) has indicated that recent research suggests that the distinction between information processing codes may not be as clear as the multiple resource theorists believe.

Rouse (1983) argues that diagnosis and compensation are two separate activities that compete with each other. The AIA model presents investigation and correction as separate stages, but the second activity may be highly dependent upon the success of the first. However, Rouse (1983) suggests that concentrating on one of the activities to the exclusion of all others may also have negative consequences. Therefore, whilst the two activities are interdependent, they have the potential for being conflicting, and Rouse asserts that this underlies the potential complexity of dealing with problem solving at multiple levels.

It is important to note that the presence of the alarm by itself may not directly suggest what course of action is required. An alarm only reports that a particular threshold has been crossed.

Monitor

Assessing the outcome of one's actions in relation to the AIAs can be presumed to be the monitor stage. It may appear to be very similar to the analyse stage in many respects, as it may involve an information search and retrieval task. Essentially, however, this mode is supposed to convey an evaluation of the effect of the corrective responses. Baber (1990) identifies three levels of feedback an operator may receive in control room tasks, these are:

- reactive;
- instrumental
- operational.

Reactive feedback may be inherent to the device, (for example, tactile feedback from a keyboard) and is characteristically immediate. Instrumental feedback relates to the lower aspects of the task, such as the typing of a command returning the corresponding message on the screen. Whereas operational feedback relates to higher aspects of the task, such as the decision to send a command which will return the information requested. These three types of feedback can be identified on a number of dimensions (Baber, 1990):

- temporal aspects;
- qualitative information content;
- relative to stage of human action cycle.

The temporal aspects refer to the relation in time for the type of feedback. Obviously reactive is first and operational is last. The content of the information relates to the degree of 'task closure' (Miller, 1968) and ultimately to a model of human action (Norman, 1986). Much of the process operator's behaviour may appear to be open-loop and therefore does not require feedback. This open-loop behaviour is due to the inherent time lag of most process systems. The literature shows that if feedback is necessary for the task, delaying the feedback can significantly impair performance (Welford, 1968). Therefore under conditions of time lag, the process operator is forced to behave in an open-loop manner. However, it is likely that they do seek confirmation that their activities have ultimately brought the situation under control, so delayed operational feedback should serve to confirm their expectations. If confirmation is sought, there is a danger that powerful expectations could lead the operator to read a 'normal' value when an 'abnormal' value is present (Moray and Rotenberg, 1989).

The operator will be receiving different types of feedback at different points in the AIAs. In the accept and correct stages they will get reactive and instrumental feedback, whereas in the monitor stage they will eventually get operational feedback. The operator is unlikely to have difficulties in interpreting and understanding reactive and instrumental feedback, if it is present, but the same is not necessarily true of operational feedback. The data presented to the operator in terms of values relating to plant items such as valves, pumps, heaters, etc., may be just as cryptic in the monitor stage as when they were requested in the investigative stage. Again the operator may be required to undertake some internal manipulation of this data in order to evaluate the effectiveness of his corrective actions, which may add substantially to the operator's mental workload.

The monitoring behaviour exhibited by humans is not continuous, but is characterized by intermittent sampling. As time passes, the process operator will become less certain about the state of the system. Crossman, Cooke *et al.* (1974) attempt to show this as a 'probability times penalty' function, where probability refers to the subjective likelihood of a process being out of specification and penalty refers to the consequences. This is balanced against the cost of sampling which means that attention will have to be diverted away from some other activity. They suggest that when payoff is in favour of sampling, the operator will attend to the process, and as soon as the uncertainty is reduced, attention will be turned to the other activities. However, they point out that monitoring behaviour is also likely to be influenced by other factors, such as: system dynamics, control actions, state changes, and the operator experienced memory decay. For example the processes may drift in an unpredictable way; operators might not know the precise effects of a control action; the process plant might be near its operational thresholds; more experienced operators might typically sample less frequently than novices, and if the operators forget values or states they might need to resample data. Crossman, Cooke *et al.* (1974) conclude from their studies that to support human monitoring of automatic systems, the system design should incorporate: a need for minimal sampling, a form of guiding the operator's activities to minimize workload, and enhanced display design to optimize upon limited attentional resources.

Conclusions

Activity in the control room may be coarsely divided into two types: routine and incident. This chapter has only considered the alarm handling aspects of the task, which have been shown to cover both routine and incident activities. However, the incident handling activities take only a small part of the operator's time, approximately 10 per cent (Baber, 1990; Rienhartz and Rienhartz, 1989) and yet they are arguably the most important part of the task. A generic structure of the task would be:

- information search and retrieval;
- data manipulation;
- control actions,
  (from: Baber, 1990)

This highlights the need to present the information to the operator in a manner that always aids these activities. Firstly, the relevant information needs to be made available to the operator to reduce the search task. The presence of too much information may be as detrimental to task performance as too little. Secondly, the information should be presented in a form that reduces the amount of internal manipulation the operator is required to do. Finally, the corrective action the operator is required to take should become apparent from both the second activity and the control interface, i.e. they can convert intention into action with the minimum of interference. It seems likely that the requirements from the alarm system may be different in each of the six stages. For example:

- conspicuity is required in the observation stage;
- time to identify and acknowledge is required in the acceptance stage;
- information to classify with related context is required in the analysis stage;
- underlying cause(s) need to be highlighted in the investigation stage;
- appropriate corrective action afforded is required in the correction stage; and
- operational feedback is required in the monitoring stage.

Therefore, it appears that alarm information should be designed specifically to support each of the stages in the alarm initiated activities (AIA) model. The difficulty arises from the conflicting nature of the stages in the model, and the true nature of alarms in control rooms, i.e. they are not single events occurring independently of each other but they are related, context-dependent and part of a larger information system. Adding to this difficulty is the range of individual differences exhibited by operators (Marshall and Shepherd, 1977) and there may be many paths to success (Gilmore, Gertman *et al.,* 1989). Therefore, a flexible information presentation system would seem to hold promise for this type of environment.

The model of AIAs (Figure 6.1) is proposed as a framework for research and development. Each of the possible alarm media has inherent qualities that make it possible to propose the particular stage of the AIA it is most suited to support. Therefore, it is suggested that speech favours semantic classification, text lists favour temporal tasks, mimics favour spatial tasks, annunciators favour pattern matching tasks and tones favour attraction and simple classification. Obviously a combination of types of information presentation could support a wider range of AIAs, such as tones and text together. These are only working hypotheses at present and more research needs to be undertaken in the AIAs to arrive at preliminary conclusions. It is proposed that:

1. the ‘observe’ stage could benefit from research in detection and applied vigilance;
2. ‘accept’ could benefit from work on group versus single acknowledgement;
3. ‘analyse’ could benefit from work on classification and decision making;
4. ‘investigate’ requires work from problem solving and diagnosis;
5. ‘correct’ needs work on affordance and compatibility; and
6. ‘monitor’ needs work on operational feedback.

However, it is already proposed that the best method of presenting alarm information will be dependent upon what the operator is required to do with the information and on the stage of AIA model the information is used. Therefore the alarm types need to be considered in terms of the AIA. This may be undertaken through a systematic comparison of combinations of alarm message across task types to investigate empirically the effect of messages type and content on performance.

In summary, it is proposed that the alarm system should support the AIA. Observation may be supported by drawing the operators' attention, but not at the expense of more important activities. Acceptance may be supported by allowing the operators to see which alarm they have accepted. Analysis may be supported by indicating to the operators what they should do next. Investigation may be supported by aiding the operators in choosing an appropriate strategy. Correction may be supported through compatibility between the task and the response. Finally, monitoring may be supported by the provision of operational feedback. The design of alarm information needs to reflect AIA, because the purpose of an alarm should not be to shock operators into acting, but to get them to act in the right way.

References

Baber, C., 1990, 'The human factors of automatic speech recognition in control rooms, unpublished PhD thesis, Aston University, Birmingham.

Baber, C., 1991a, *Speech technology in control room systems: a human factors perspective,* Chichester: Ellis Horwood.

Baber, C., 1991b, Why is speech synthesis inappropriate for control room applications? In Lovesey, E.J. (Ed.) *Contemporary Ergonomics: Ergonomics Design for Performance,* London: Taylor & Francis.

Baber, C., Stammers, R.B. and Taylor, R.T., 1990, Feedback requirements for automatic speech recognition in control room systems, in Diaper, D., Gilmore, D., Cockton, G. and Shackel, B. (Eds) *Human-Computer Interaction: INTERACT '90,* pp. 761–6, Amsterdam: North-Holland.

Bainbridge, L., 1983, The ironies of automation, *Automatica,* **19** (6), 775–9.

Bainbridge, L., 1984, Diagnostic skill in process operation, *International Conference on Occupational Ergonomics,* 7–9 May, Toronto.

Boreham, N.C., 1985, Transfer of training in the generation of diagnostic hypotheses: the effect of lowering fidelity of simulation, *British Journal of Education Psychology,* **55,** 213–23.

Brehmer, B., 1987, Models of diagnostic judgements, in Rasmussen, J., Duncan, K. and Leplat, J. (Eds) *New Technology & Human Error,* Chichester: Wiley.

Broadbent, D., 1990, Modelling complex thinking, *The Psychologist,* **3** (2), 56.

Coury, B.G. and Pietras, C.M., 1989, Alphanumeric and graphical displays for dynamic process monitoring and control, *Ergonomics,* **32** (11), 1373–89.

Crossman, E.R.F.W. and Cooke, J.E., 1974, Manual control of slow response systems, in Edwards, E. and Lees, P.P. (Eds) *The Human Operator in Process Control,* London: Taylor & Francis.

Crossman, E.R.F.W., Cooke, J.E. and Beishon, R.J., 1974, Visual attention and the sampling of displayed information in process control, in Edwards, E. and Lees, P.P. (Eds) *The Human Operator in Process Control.* London: Taylor & Francis.

Duff, S.C., 1989, Reduction of action uncertainty in process control systems: the role of device knowledge, in *Contemporary Ergonomics 1989,* Proceedings of the Ergonomics Society 1989 Annual Conference; 3–7 April, London: Taylor & Francis.

Easterby, R., 1984, Tasks, processes and display design, in Easterby, R. and Zwaga, H. (Eds) *Information Design,* Chichester: Wiley.

Edwards, E. and Lees, P.P., 1974, *The Human Operator in Process Control,* London: Taylor & Francis.

Edworthy, J. and Loxley, S., 1990, Auditory warning design: the ergonomics of perceived urgency, in Lovesey, E.J. (Ed.) *Contemporary Ergonomics 1990: Ergonomics setting the standards for the 90s;* 3–6 April, pp. 384–8, London: Taylor & Francis.

Elkind, J.I., Card, S.K., Hochberg, J. and Huey, B.M., 1990, (Eds) *Human Performance Models for Computer-Aided Engineering,* Boston: Academic Press;

Eysenck, M.W., 1984, *A Handbook of Cognitive Psychology,* London: Lawrence Erlbaum Associates.

Farah, M.J., 1989, Knowledge from text and pictures: a neuropsychological perspective, in Mandl, H. and Levin, J.R. (Eds) *Knowledge Aquisition from Text and Pictures,* pp. 59–71, North Holland: Elsevier.

Gilmore, W.E., Gertman, D.I. and Blackman, H.S., 1989, *User-Computer Interface in Process Control,* Boston: Academic Press.

Goodstein, L.P., 1985, *Functional Alarming and Information Retrieval,* Denmark: Risø National Laboratory, August Risø-M-2511. 18.

Goodstein, L.P., Andersen, H.B. and Olsen, S.E., 1988, *Tasks, Errors and Mental Models*, London: Taylor & Francis.

Hale, A.R. and Glendon, A.I., 1987, *Individual Behaviour in the Control of Danger,* Amsterdam: Elsevier.

Hoc, J-M., 1988, *Cognitive Psychology of Planning,* London: Academic Press.

Hockey, G.R.J., Briner, R.B., Tattersall, A.J. and Wiethoff, M., 1989, Assessing the impact of computer workload on operator stress: the role of system controllability, *Ergonomics,* **32** (11), 1401–18.

Jenkinson, J., 1985, *Alarm System Design Guidelines,* Central Electricity Generating Board, September, GDCD/CIDOCS 0625.

Johannsen, G., 1988, Categories of human operator behaviour in fault management situations, in Goodstein, L.P., Andersen, H.B. and Olsen, S.E. (Eds) *Tasks, Errors and Mental Models,* pp. 251–58, London: Taylor & Francis.

Kortlandt, D. and Kragt, H., 1980, Process alarm systems as a monitoring tool for the operator, in *Proceedings of the 3rd International Symposium on Loss Prevention and Safety Promotion in the Process Industries.;* September 15–19, Basle, Switzerland, Vol. 1, pp. 10/804–10/814.

Kragt, H. and Bonten, J., 1983, Evaluation of a conventional process alarm system in a fertilizer plant. *IEEE Transactions on Systems, Man and Cybernetics,* **13** (4), 589–600.

Landeweerd, J.A., 1979, Internal representation of a process, fault diagnosis and fault correction, *Ergonomics,* **22** (12), 1343–51.

Mann, T.L. and Hammer, J.M., 1986, Analysis of user procedural compliance in controlling a simulated process, *IEEE Transactions on Systems, Man and Cybernetics;* **16** (4), 505–10.

Marshall, E. and Shepherd, A., 1977, Strategies adopted by operators when diagnosing plant failures from a simulated control panel, in *Human Operators and Simulation,* pp. 59–65. London: Institute of Measurement & Control.

Miller, R.B., 1968, Response time in man-Computer conversational Transactions, *Proceedings of the Spring Joint Computer Conference,* **33**; 409–21, Reston, Virginia: AFIRS Press.

Molitor, S., Ballstaedt, S-P. and Mandl, H., 1989, Problems in knowledge acquisition from text and pictures, in Mandl, H. and Levin, J.R. (Eds) *Knowledge Acquisition from Text and Pictures,* pp. 3–35, North Holland: Elsevier.

Moran, T.P., 1981, An applied psychology of the user, *Computer Surveys,* 13 (1), 1–11.

Moray, N., 1980, The role of attention in the detection of errors and the diagnosis of failures in man-machine systems, in Rasmussen, J. and Rouse, W.B. (Eds) *Human Detection and Diagnosis of System Failures,* New York: Plenum Press.

Moray, N. and Rotenberg, I., 1989, Fault management in process control: eye movements and action, *Ergonomics;* **32** (11), 1319–42.

Morris, N.M. and Rouse, W.B., 1985, The effects of type of knowledge upon human problem solving in a process control task, *IEEE Transactions on Systems, Man and Cybernetics,* **15** (6), 698–707.

Newell, A. and Simon, H.A., 1972, *Human Problem Solving,* Englewood Cliffs, NJ: Prentice Hall.

Norman, D.A., 1986, Cognitive engineering, in Norman, D.A. and Draper, S.W. (Eds) *User Centred System Design, Hillsdale,* NJ: Lawrence Erlbaum Associates.

Pal, J.K. and Purkayastha, P., 1985, Advanced man-machine interface design for a petroleum refinery plant, in Johannsen, G., Mancini, G. and Martensson, L. (Eds) *Analysis, Design and Evaluation of Man-Machine Systems,* pp. 331–37, Italy: Commission of the European Communities.

Perrow, C., 1984, *Normal accidents: Living with high risk technology,* New York: Basic Books.

Pew, R.W. and Baron, S., 1982, Perspectives on human performance modelling, in Johannsen, G. and Rijnsdorp, J.E. (Eds) *Analysis, Design & Evaluation of Man-Machine Systems,* Duesseldorf: IFAC.

Pew, R.W., Miller, D.C. and Freehrer, C.E., 1982, Evaluating nuclear control room improvements through analysis of critical operator decisions, *Proceedings of the Human Factors Society 25th Annual Meeting,* pp. 100–4.

Rasmussen, J., 1976, Outlines of a hybrid model of the process plant operator, in Sheriden, T.B. and Johannsen, G. (Eds) *Monitoring Behaviour and Supervisory Control,* New York: Plenum Press.

Rasmussen, J., 1983, Skills, rules and knowledge; signals, signs and symbols, and other distinctions in human performance models, *IEEE Transactions on Systems, Man and Cybernetics,* **13** (3).

Rasmussen, J., 1984, Strategies for state identification and diagnosis in supervisory control tasks, and design of computer based support systems, in Rouse, W.B. (Ed.) *Advances in Man-Machine Systems Research,* pp. 139–93.

Rasmussen, J., 1986, Information processing and human-machine interaction, *An Approach to Cognitive Engineering,* North-Holland: Amsterdam.

Reason, J., 1988a, The Chernobyl errors, *Bulletin of the British Psychological Society,* **40**, 201–6.

Reason, J., 1988b, Framework models of human performance and error: a consumer guide, in Goodstein, L.P., Andersen, H.B. and Olsen, S.E. (Eds) *Tasks, Errors and Mental Models,* London: Taylor & Francis.

Reason, J., 1988c, Cognitive aids in process environments: prostheses or tools? In Hollnagel, E., Mancini, G. and Woods, D.D. (Eds) *Cognitive Engineering in Complex Dynamic Worlds,* pp. 7–14.

Reason, J., 1990, *Human Error* Cambridge: Cambridge University Press.

Reed, J. and Kirwan, B., 1991, An assessment of alarm handling operations in a central control room, in Quéinnec, Y. and Daniello, F. (Eds) *Designing for Everyone,* London: Taylor & Francis.

Reinartz, S.J. and Reinartz, G., 1989, Analysis of team behaviour during simulated nuclear power plant incidents, In Megaw, E.D. (Ed.) *Contemporary Ergonomics 1989,* Proceedings of the Ergonomics Society 1989 Annual Conference 3–7 April, pp. 188–93, London: Taylor & Francis.

Rouse, W.B., 1983, Models of human problem solving, *Automatica,* **19,** 613–25.

Rouse, S.H. and Rouse, W.B., 1982, Cognitive style as a correlate of human performance in fault diagnosis tasks, *IEEE* Transactions on Systems, Man and Cybernetics, **12** (5), 649–52.

Singleton, W.T., 1989, *The Mind at Work,* Cambridge: Cambridge University Press.
Sorkin, R.D., 1989, Why are people turning off our alarms? *Human Factors Bulletin,* **32**, 3–4.
Stanton, N.A., 1992, 'Human factors aspects of alarms in human supervisory control tasks,' unpublished Phd thesis, Aston University: Birmingham.
Stanton, N.A., 1993, Operators reactions to alarms: fundamental similarities and situational differences, *Proceedings of the Conference on Human Factors in Nuclear Safety,* Le Meridien Hotel, London, 22–23 April.
Stanton, N.A. and Booth, R.T., 1990, The psychology of alarms. In Lovesey, E.J. (Ed.) *Contemporary Ergonomics,* London: Taylor & Francis.
Stanton, N.A., Booth, R.T. and Stammers, R.B., 1992, Alarms in human supervisory control: a human factors perspective, *International Journal of Computer Integrated Manufacturing,* **5** (2), 81–93.
Su, Y-L. and Govindaraj, T., 1986, Fault diagnosis in a large dynamic system: experiments on a training simulator, *IEEE Transactions on Systems, Man and Cybernetics,* **16** (1), 129–41.
Swain, A.D. and Weston, L.M., 1988, An approach to the diagnosis and misdiagnosis of abnormal conditions in post-accident sequences in complex man-machine systems, in Goodstein, L.P., Andersen, H.B. and Olsen, S.E. (Eds) *Tasks, Errors and Mental Models,* London: Taylor & Francis.
Welford, A.T., 1968, *Fundementals of Skill,* London: Methuen.
Wickens, C.D., 1984, *Engineering Psychology and Human Performance,* Columbus, Ohio: Merrill.
Wickens, C.D. and Kessel, C, 1981, Failure detection in dynamic systems, in Rasmussen J. and Rouse, W.B. (Eds) *Human Detection and Diagnosis of Systems Failures,* New York: Plenum Press, pp. 155–69.
Woods, D.D., 1988, Coping with complexity: the psychology of human behaviour in complex systems, in Goodstein, L.P., Andersen, H.B. and Olsen, S.E. (Eds) *Tasks, Errors and Mental Models,* pp. 128–48, London: Taylor & Francis.