AIRCRAFT COMPUTERS. See AIR TRAFFIC.
AIRCRAFT. See AIR TRAFFIC.

ALARM SYSTEMS

The role of personnel in most human–machine systems involves a number of generic tasks, such as monitoring system performance, detecting disturbances or off-normal conditions, assessing the system's status, planning and implementing responses to disturbances, and obtaining feedback regarding response effectiveness. To support these tasks, personnel obtain information and execute control using human–machine interfaces, such as alarms, displays, and controls. (Although the personnel using alarm systems include operations, maintenance, or other functions, the term "operators" is used for simplicity.)

In many complex human–machine systems, monitoring and detection functions easily become overwhelming because of the large number of individual parameters and conditions involved. Therefore, support is generally provided for these activities by an alarm system. The alarm system is one of the primary means by which abnormalities and failures come to the attention of the operators.

In this context, an "alarm system" is actually an automated monitoring system that alerts operators via visual and/or auditory displays when parameters deviate from specified limits, called set points. The basic alarm system functions include

- Alerting the operator to the fact that a system or parametric deviation exists
- Informing the operator about the priority and the nature of the deviation
- Guiding the operator's initial response to the deviation
- Confirming whether the operator's response corrected the deviation

Alarm systems consist of sensing, processing, and display hardware and software. In the broadest sense, an alarm reflects a parameter, component, system, or function that is in a state requiring the attention of personnel. The state itself may or may not be adverse. In a narrow sense, the term alarm means an attention-eliciting message given to personnel regarding a potentially adverse deviation of a parameter, component, system, or function from its expected performance.

Although alarm systems play an important role in system operations, they have often posed challenges to the personnel who must rely on them. These challenges often make alarm systems difficult to use when they are needed the most—during significant system disturbances. The most common alarm design deficiencies are

1. too many alarms
2. too many false alarms
3. difficulty distinguishing alarms from normal status indications
4. poor alarm organization (which makes it difficult to see the connections between indidual alarms and between alarms and components

5. poor location (e.g., alarms not in the operator's direct view and not located near associated controls and displays
6. insufficient salience coding (i.e., important alarms fail to draw the operator's attention)
7. inadequate message design (e.g., poor labeling, poor legibility, ambiguous messages)
8. poor acoustic design (e.g., masking of alarms and irritating or distracting warnings)
9. inadequate identification of the absolute or relative importance among alarms.

To understand the reasons for these problems, it is important to consider the role of alarms in human information processing.

OPERATOR USE OF ALARM INFORMATION

The primary function of an alarm system is to support the detection of off-normal situations. Human fault detection can be described in terms of signal detection theory (SDT) (1). Within this framework, the operator and the alarm system constitute an alerted-monitor system (2,3). The automated monitor is the alarm system, which monitors parameters to detect off-normal conditions. When a parameter exceeds the criterion of the automated monitor, the human monitor is alerted and then must detect, analyze, and interpret the signal as a false alarm or a true indication of an off-normal condition. The human monitor also assesses system parameters independently of the alarm system. Both the human and automated monitors have their own decision criteria and sensitivity. The decision criterion refers to the amount of evidence that is needed before a conclusion is made that a signaled event is actually present (this is sometimes called response bias). Sensitivity is the resolution of the system that determines the ease with which true disturbances are distinguished from routine fluctuations.

SDT research has many implications for understanding how operators process alarm information. First, the response criterion is affected by context, that is, the expected probability that an event will occur and the payoff structure (rewards and penalties for making correct and incorrect detections, respectively). Significant off-normal events in many systems may have a low probability of occurring. Therefore, operators may be disinclined to decide that they have actually occurred. There is a conflict between the cost to productivity in responding (unnecessarily) to a false alarm versus the potentially significant cost of failing to respond to a true alarm. When disturbances have a low probability, operators rely on redundant and supplemental information to confirm the alarmed condition. Upon verification of several confirmatory indicators, the operator accepts the alarm information as indicating an actual off-normal condition (compared with a spurious condition).

Once operators conclude that the alarm information represents a valid system disturbance, they actively try to construct a coherent, logical explanation to account for their observations. This cognitive activity is called situation assessment and involves two related concepts: the situation model and the mental model. Operators develop and update a mental representation of the factors known or hypothesized to be affecting the system at a given time. The mental representation resulting from situation assessment, called a *situation model,* is the operator's understanding of the specific current situation and is constantly updated as new information is received. To construct a situation model, operators use their general knowledge and understanding about the system and how it functions to interpret what they observe and to understand its implications. This general knowledge, commonly called the operator's *mental model,* constitutes an internal representation of the physical and functional characteristics of the system and its operation.

Based on the situation model, the operator plans responses to the disturbance that are designed to bring the system back to a safe state. The plans are carried out, and the operator looks for feedback to indicate the success of the plans. Alarms play a role throughout this process. When the alarm system poses problems, operator performance may suffer.

The problems encountered by the operator in using alarm systems are illustrated by examining circumstances when systems deviate from normal operating conditions. The nuclear power plant incident at Three-Mile Island (TMI) is a good example. The President's Commission on TMI indicated that during the first few minutes of the accident, more than 100 alarms were activated in the control room (4). The operators had no assistance in distinguishing significant alarms (i.e., situations requiring operator attention) from less important signals (i.e., situations not requiring operator attention). The Rogovin report on the incident was more specific in identifying alarm system deficiencies and contributions to safety (5). The report indicated that on "the morning of the accident, the alarms were of little use due to the number that were flashing and their almost random locations." Some of the important alarms were not located in direct view of the operators. Although auditory signals were associated with these alarms, they could not be distinguished from other alarms because a single button caused the auditory signals and flashing lights to stop for all alarms. Operators indicated that the constant buzzing of auditory alarms and flashing lights were distracting and made their jobs more difficult.

The types of problems experienced by the operators at TMI are typical of the problems faced by operators of many complex human–machine systems, such as process control facilities, aircraft, and medical systems.

Alarm overload is a significant challenge to operators. They often have problems detecting and interpreting alarms. It has been shown that an operator's ability to detect off-normal events is reduced as the number of alarms presented increases (6,7,8). In fact, it has been generally found that, as demands on operators increase, fault detection capability decreases (9). When the number of alarms is large, the operator's information processing ability becomes overloaded, and performance suffers because of high workload (2). This occurs because the operator's cognitive resources are in short supply, and the resulting changes in behavior increase the probability of error. First, an operator may adopt inappropriate alarm sampling strategies which make the accurate diagnosis of system anomalies less likely (10,11). Under normal conditions, a sampling strategy based on successive observations of weakly related variables is an appropriate strategy. However, once a disturbance begins, a more appropriate strategy is to sample correlated variables because this facilitates detecting and recognizing a system/component failure. The "normal" sampling strategy plus the operator's low expectancy of problems can delay realization of an event. Alarm system design features, for example, grouping alarms by system and function, help minimize this problem. This type of organization enables operators to perceptually integrate the alarms into high-level information, for example, perceiving several low-level alarms as a higher level event, such as "Pump A is tripped."

A second result of alarm overload is a shift toward simpler information-processing strategies. Reason has indicated that operators under high workload conditions employ two problem-solving heuristics: "similarity matching" and "frequency gambling" (12,13). These strategies give rise to a number of "basic error tendencies" in human performance which account for many human errors. Similarity matching reflects the tendency for the operator to attempt to match a perceived information pattern (such as a pattern of alarm signals) with known, remembered patterns. When the perceived information partially matches more than one mental representation, the discrepancy is resolved by selecting the one most frequently experienced. This is the "frequency gambling" heuristic. There are alarm system features that assist the operator in processing information and thereby help prevent errors. An example of such a feature is the generation of alarms when (1) "unexpected" alarms (based on the current pattern) occur and (2) an "expected" alarm (based on the current pattern) does not occur. Such alarms call the operator's attention to "outlier" alarms which are likely to be missed because of a tendency to focus on indications that confirm the operator's current representation. The use of alarm overview displays and hierarchal alarm displays also help minimize these types of errors.

In addition to alarm overload, nuisance alarms are another significant problem with alarm systems. As discussed previously, fault detection performance is a function of the entire alerted-monitor system. Optimizing the signal detection parameters for one component of the system may not optimize performance of the entire two-stage system (2). Thus, when the response criterion of the alarm system is set to maximize the number of disturbances detected, the number of false alarms increases. This problem occurs when alarm parametric setpoints are too close to the normal operating value or too close to the normal value drift. Although this may provide an early alert to a potential system disturbance, many false alarms are created because of momentary fluctuations in parametric values. When there are many false alarms, an operator may lose confidence in the system and adopt a more conservative criterion. Such interactions between automated and human monitors can result in poor overall performance.

The need to address these deficiencies has led to advances in nearly every aspect of alarm system design. For example, alarm processing has been provided to facilitate identification of critical alarms, and alarm displays have been improved to facilitate association of alarms and other system information. In addition, many alarm systems are designed with alarm management facilities enabling personnel to sort alarms along dimensions, such as time and system, and to interrogate the alarm system to obtain detailed information about specific alarms of interest.

ALARM SYSTEM DESIGN

This section addresses alarm system design characteristics relative to the alarm functions identified earlier. The specific characteristics are alarm selection, alarm states, alarm prioritization and processing, alarm availability, display, control, automated and modifiable characteristics, reliability and maintenance, location and integration of alarms into the workplace, and alarm response procedures. The main discussion pertains to general alarm-design considerations that apply to complex systems, such as process control facilities, aircraft and surface transportation systems, and medical devices. Alarm systems for simpler systems may not involve all of the considerations presented.

Alarm Selection

The first consideration in alarm system design is selecting conditions to be monitored. Important considerations include (1) the categories (i.e., the events or system states) from which alarms are identified, (2) the criteria to be used in selecting alarm parameters to represent the categories, and (3) criteria for determining set points.

The situational categories from which alarms are selected are personnel hazards, safety, productivity, and investment protection (indicators of damage to important and costly equipment). Such conditions often require operator attention and action to prevent a system disturbance from resulting in negative consequences. For purposes of this chapter, a distinction is made between alarms and annunciators, although the two terms are often used interchangeably in the literature. An annunciator is a status indicator that does not necessarily require the operator's attention. Annunciators should not be presented by the alarm system displays because they distract operators from attending to actual alarms.

Once the parameters representing these conditions are selected, set points are determined. A set point is the value of a monitored parameter which defines the boundary between the parameter's normal range and an alarm condition. An alarm condition exists when the parameter exceeds the normal range defined by the upper and/or lower set points. Graded alarms may have multiple setpoints outside of the normal range that produce alarms indicating increasing levels of severity of an abnormal condition (e.g., low level, low-low level, etc.).

Set points need to be carefully selected to ensure that operators monitor and take appropriate action in a timely manner. To achieve this, set points are specified at levels that are sufficiently different from the actual limits to allow sufficient time to respond. However, considering the signal detection issues discussed before, determining alarm set points should consider the trade-off between timely alerts to off-normal conditions and the creating of nuisance alarms by establishing set points so close to the "normal" operating values that occasional excursions of no real consequence are expected.

Alarm States

Alarms have four states: inactive, new, acknowledged, and cleared. When the parameter is in the normal range, the alarm is inactive. A new alarm occurs when a monitored parameter exceeds a specified limit (set point). The deviation is evaluated by the processing portion of the alarm system, and an indication is conveyed to the operator through the display portion of the alarm system. The alarm remains new until it is acknowledged.

An alarm is in the acknowledged state when the operator has provided some type of input to the alarm system (such as pressing a button) to indicate receipt of the alert or message provided by the alarm system. The act of acknowledging an alarm typically causes the attention-getting characteristics of the alarm display to cease or decrease (e.g., the auditory tone stops, and the flashing display changes to a steady illumination).

An alarm is cleared when the alarmed parameter returns from an alarmed state to its normal range. Some alarm systems provide indications when the parameter enters the normal range. The operator may be required to acknowledge the alarm to "clear" it.

Alarm Prioritization and Processing

Not all alarms are of equal importance. Therefore, prioritizing alarms is helpful to operators. Alarm priority is typically based on the immediacy of required operator action and challenges to safety. The prioritization scheme should be logical so that those alarms of the highest safety significance receive the highest priority and so that the prioritization appears reasonable to operators. It is best to limit the number of priority levels to four or less. Prioritization schemes with many levels require operators to devote excessive attention to the priority level and thus reduce the benefits of prioritization.

Priorities are determined by engineering analyses during system design. This type of a priori analysis identifies the static priority of alarmed conditions. Static in this case means that the relative importance of alarms is always the same and does not change as a result of real-time analyses during system operation. More sophisticated systems use alarm processing techniques to dynamically prioritize alarms on the basis of the existing situation and system condition.

Alarm processing techniques were developed to support operators in coping with the high volume of alarms that occur during major system disturbances by reducing the number of alarms. Two general classes of alarm processing techniques are discussed: signal validation and condition processing. When instrumentation failures (such as a failed sensor) occur, biased or false signals are generated. The use of these signals by the alarm system may result in the presentation of either false or nuisance alarm messages. Such alarm messages are misleading and may interfere with the operator's situation assessment or reduce the crew's confidence in future alarm messages. Signal validation is a set of techniques by which signals from redundant or functionally related sensors are automatically evaluated to determine whether a true alarm condition exists, thereby avoiding presentation of spurious alarms to operators. Such techniques include the analysis of normal signal drift and noise signals to eliminate those that momentarily exceed the set points but do not indicate a true alarm condition. Alarm conditions not eliminated by the alarm signal processing may be evaluated further by alarm condition processing.

Alarm condition processing is the rules or algorithms used to determine the operational importance and relevance of alarm conditions. This is done to determine how the alarm messages associated with these conditions should be presented to the operator. For the purposes of this discussion, four classes of processing techniques are defined: nuisance alarm processing, redundant alarm processing, significance processing, and alarm generation processing. Each processing technique changes the resulting information provided to operators.

Nuisance alarm processing includes techniques that seek to eliminate alarms which have no operational safety importance. Some examples of this class of techniques are time-delay and mode-dependent processing. Time-delay processing applies time averaging and/or time delay to the alarm inputs to allow filtering noise signals and to eliminate unneeded momentary alarms (such as those resulting from instrument noise). Mode-dependence processing evaluates alarm signals for their relevance to the current system mode or configuration. If a component's parametric value represents a fault in some system configurations and not others, it should be alarmed only in the appropriate configurations. For example, the fact that a particular pump has a low discharge pressure may only indicate a fault when the associated fluid system is configured to perform a particular function. Other discharge pressures may be appropriate when the fluid system is configured to perform different functions. Mode processing would allow the alarm message to be presented only when it is meaningful to the fluid system's current configuration.

Redundant alarm processing includes techniques that analyze for alarm conditions that are valid but are considered less important because they provide information that is redundant to other alarms and logically provide no new information. For example, a single event (cause) invariably leads to subsequent alarmed events that are the direct consequence of the initial event. Processing is used to eliminate alarms that follow as logical consequences. Such techniques may, however, minimize information used by the operator for (1) confirmation that the situation represented by the "true" alarm has occurred, (2) situation assessment, and (3) decision-making. Thus, in addition to quantitatively reducing alarms, processing methods may qualitatively affect the information presented to the operating crew.

Significance processing includes techniques that analyze for alarm conditions that are valid but are considered less important because of their significance compared with other alarm conditions. For example, alarms for a condition representing a threat to personnel safety may be displayed whereas a concurrent alarm associated with minor equipment failure is not presented.

Alarm generation processing includes techniques that analyze the existing alarm conditions and, based on the evaluation, generate alarm messages which (1) give the operator higher level or aggregate information, (2) notify the operator when "unexpected" alarm conditions occur, and (3) notify the operator when "expected" alarm conditions do not occur. The generation of alarm conditions and their resulting alarm messages presents an interesting paradox. Alarm systems should facilitate the reduction of errors which often reflect the overloaded operator's incomplete processing of information. Alarm generation features may help mitigate these problems by calling the operator's attention to plant conditions that are likely to be missed. The single most significant problem with alarms systems, however, is the high number of alarm messages. Because alarm generation creates additional messages, it may potentially exacerbate the problem.

There are two additional aspects to alarm processing to be considered in alarm system design: degree of alarm reduction and complexity of processing. The relationship between the degree of alarm reduction resulting from processing and its impact on operator performance is not fully understood, that is, although we know that most systems have too many alarms, we do not know how much reduction is necessary to improve the situation meaningfully.

The processing methods applied should not be so complex that operators have difficulty evaluating the meaning or validity of the resulting alarm messages. To support the understandability of alarm processing systems, inputs to the alarm processing system (e.g., sensor data) should be available to the operators. This data may be needed, for example, if the pattern of alarm messages appears to be contradictory or if operators suspect that there is a problem with the processing system, such that the results of alarm processing are incorrect.

Alarm Availability

Based on the types of analyses previously discussed, some alarms are considered more important than others. Alarm availability is the method by which the differences in alarm importance are used to determine which alarms are made available to operators. Three techniques have been used: filtering, suppression, and priority coding. Note that the terms "filtering" and "suppression" are often used interchangeably in the literature, although different meanings are identified here.

Alarm filtering is the process of eliminating alarms determined to be less important. Filtered alarms are not available to operators. Filtering should be employed only where alarms have no operational significance. Thus, only alarms that have no operational significance to operators should be filtered.

In alarm suppression, less important alarms are not presented to the operators on the primary alarm displays but are available on auxiliary displays (which the operators can retrieve) or are presented automatically when they become high in importance (e.g., when more important alarm conditions have cleared).

The third option is dynamic priority coding where all alarms are presented to operators but information about their priority is given. For example, more important alarms may be color coded red whereas less important alarms are white.

Although a specific alarm system employs a combination of these approaches, there are clear trade-offs among them. Filtering eliminates the possibility that unimportant alarms distract the operators. However, the designer may be removing information that may be used by operators for other purposes. In addition, the designer must be certain that the processing method is adequately validated and functions appropriately in all conditions. Suppression provides the potential benefits of filtering by removing potentially distracting alarms. But, because the suppressed alarms are accessible, they potentially impose an additional workload by requiring operator action to retrieve them. Priority coding, on the other hand, does not conceal any information from operators because all information is presented. However, the operator is required to search for the important, high-priority alarms among the potentially distracting visual clutter of less important alarms.

Alarm Display

In older systems, alarm displays are relatively simple. Each alarmed parameter has an indicator, for example, a small plastic rectangular "tile" with a label briefly indicating what the alarm is. The individual alarm tiles are grouped together into alarm matrices. When an alarm set point is reached, an auditory signal is the first indication. Upon hearing the auditory signal, the operator looks at an alarm panel. At the same time as the auditory alert, the specific alarm tile begins to flash. So finding the alarm is relatively easy, as long as not too many alarms come in within a short time. In some systems, alarm tiles are augmented by alarm message printers that print out a chronological list of alarms as they come in. In these older systems, the alarm system is typically an information system separate from the other information systems. The alarm system alerts operators to off-normal conditions. Then the operators consult other indicators for specific information.

There are a number of general types of alarm displays.

- Spatially dedicated continuously visible (SDCV) alarm displays—A spatially dedicated alarm is always displayed in the same physical location. Continuously visible means that a parallel presentation method is used, that is, the alarm information is always displayed to the operator, as opposed to serial presentation methods in which the operator must select the information to be seen. The tile display, described before, is an example of an SDCV alarm. More recently, tile-like video displays have been used where operators are seated at a workstation.
- Temporary alarm displays—Alarm message lists are a typical implementation of a temporary alarm display. Messages appear only when the alarm is activated. Depending on the design, temporary alarms may or may not appear in spatially dedicated locations.
- Integrated alarms—Alarm information is presented as an integral part of other displays, such as process displays. For example, if alarms are built into a system mimic display, trouble with a component, such as a pump, is depicted by a change in color or flashing of the pump icon. These types of displays may have fixed or variable locations.

SDCV alarm displays are superior to temporary alarm displays during high-alarm conditions. The fixed locations are often thought to provide perceptual advantages of rapid detection because operators know alarms by position. Thus when an alarm flashes, they often do not have to read its label in detail. Further, operators begin to recognize the patterns of alarms associated with familiar disturbances. Although message lists typically provide more information than alarm tiles, they are problematic because of the workload associated with reading individual messages. Because spatial cues to the nature of the alarms are not available, operators must read each alarm as it comes in to know what it is.

In recent years, alarm displays have become considerably more complex. Although audio and visual components are still important attention-directing features of alarm displays, alarms are becoming increasingly integrated into normal information displays used for system monitoring and decision making. One of the reasons for this trend is the recognition that cognitive processing is facilitated by integrating information into a single object (14) or display (15). It is thought that integrated displays enhance the ability of the operators to process information in parallel, thus lowering cognitive workload. Data integration also enables operators to understand the relationships between display elements better and to assess the situation more rapidly and accurately.

Because of their superiority in high-alarm conditions, SDCV alarm displays should be considered for alarms that address (1) safety and mission-critical parameters, (2) situations that require short-term response by the operators, (3) the main alarms used by operators in diagnosing and responding to plant upsets, and (4) the principal alarms used by operators to maintain an overview of plant and system status.

Some of the general design considerations for alarm displays follow. The discussion is divided into the following topics: display of alarm states, message design, coding, shared alarms, and organization.

Display of Alarm States. Each alarm state (i.e., new, acknowledged, cleared, and inactive) should have a unique presentation to support the operator's ability to rapidly distinguish among them. New alarms are indicated both visually (e.g., flashing) and audibly. After the operator has acknowledged an alarm (e.g., pressed the acknowledge button), the alarm display changes to a visually distinct acknowledged state, and the alerting function (e.g., audible tone) ceases.

If the operator is required to take action when an alarm clears, the return to normal conditions is also indicated visually and by audible means, called ringback. This is not needed for all alarms but is useful when it is important for the operator to know immediately when the deviation has cleared or when the deviation is not expected to clear for some time. Techniques that are employed include a special flash rate (one-half the normal flash rate is preferred to allow discrimination), reduced brightness, or a special color that is consistent with the overall control room color-coding scheme. Cleared alarms should have a dedicated, distinctive audible signal of finite and relatively short duration.

Inactive alarms are best indicated by an absence of visual and auditory alarm features. This practice is referred to as the dark board (or blackboard) concept of alarm display because it results in a dark display medium (not illuminated) when all monitored plant parameters are in the normal range. Under such circumstances, a new alarm is easily identified.

Message Design. Alarm messages include information, such as

- Alarm title
- Time of occurrence
- Alarm source, that is, the particular sensor or group of sensors supplying the signal
- Alarm priority
- Set point and parametric values
- Required immediate operator actions
- Reference to procedure for more detailed follow-up actions

The extent to which all of this information is included in an alarm message is constrained by the type of display system available. More information is provided when alarms are presented on a video display unit (VDU). A tile-based system necessarily displays less information, requiring that detailed information be provided through other alarm displays, such as supplemental VDUs and printers. It is important that the format of messages on alarms be consistent for all alarms and consistent across types of alarms and between VDU and printed message displays.

Alarm title or legend text should be clearly understandable, using standard terminology and addressing conditions specifically. For example, it is preferable to identify the parameters and states (e.g., HIGH PRESSURE) specifically instead of using one legend for multiple parameters or multiple states (e.g., TEMPERATURE-PRESSURE or HIGH-LOW).

Operators often want to know the chronology of alarms, especially for diagnosis. Providing the time the alarm was triggered is beneficial for these types of activities.

The display of important alarms takes precedence over lower alarms. For alarms that are not SDCV displays, sufficient display area is needed for simultaneously viewing all high-priority alarms. For example, operators should not have to page or scroll a display to view high-priority alarms. Codes are used to indicate alarm priority levels. Coding methods include color, position (top to bottom), shape, and symbolic coding. Color and position are especially effective visual coding methods. However, coding by position should not disrupt the functional grouping of alarms. Coding is discussed in greater detail later.

If an alarm condition requires verification before action is taken, then relevant set-point limits should be included in the alarm message. However, one needs to consider whether providing set points alone (and not current parametric value) leads operators to assume that the condition is near the set point when, in fact, it is well above or below it. Whenever this is an important consideration, actual parametric values may be as important as set points and should be provided.

Where practical, immediate operator actions are presented or referenced. For example, immediate operator actions are provided in alarm response procedures (discussed later) that are clearly and simply keyed to an alarm tile and located nearby for easy and quick reference.

Coding. Coding for alarm priority and state was mentioned before. Coding is the use of a system of symbols, shapes, colors or other variable sensory stimuli to represent specific information. Coding schemes facilitate rapid detection and interpretation of alarms by operators. Each level of a coding dimension must be easily and readily distinguishable from the other levels. A formal coding scheme is established and formally documented which encompasses all coding dimensions (e.g., color, shape, brightness, textures/pattern, flashing, and sound) and specifies a hierarchical order of salience. Then the coding scheme is systematically applied to alarm information so that the most important alarm information is associated with the most salient coding. Because coding adds to display complexity, it is a good practice to minimize the number of different coding techniques.

The primary coding methods used for alarms are visual and auditory. These techniques are discussed in the remainder of this section.

Color is one of the most effective types of coding. Colors should be easily discriminated, each color with a single, precise meaning consistent with its use in the rest of the system's operator interfaces. It is important that the use of color is consistent with and, if possible, takes advantage of operators' already learned color associations.

A flashing visual signal is typically associated with all important alarms to ensure that the operator's attention is captured and directed to these alarms under any and all conditions. To achieve this, flash rates from three to five flashes per second with approximately equal on and off times are effective.

For transilluminated displays, the brightest state should be no more than 300% brighter than the inactivated state (but not annoy operators) and the dim state at least 10% brighter than the inactivated state. For VDU displays, the bright state should be at least 100% brighter than the inactivated state. VDUs can be used to display two brightness levels, whereas transilluminated alarms may display more than two levels of brightness. It is ineffective to use low-intensity indications (e.g., dark red) in the periphery of the visual field where color coding is used because they may not be readily detected.

Under high alarm-volume conditions, consider suppressing or delaying the alerting indications (e.g., visual flashing) for those alarm conditions that (1) do not require immediate response and (2) do not indicate a challenge to safety. In addition, redundant codes (e.g., fast flashing or bright illumination) are used to assist operators in detecting the more significant alarm messages and to reduce distraction from less important alarms.

Auditory signals draw attention to alarms even when operators are not attending to the visual alarm displays. Consequently, it is common practice to provide auditory cues for all new alarms under normal operating conditions. The number and placement of sound sources should be such that auditory signals are free of distortion and are equally audible at any operator work station in the primary operating area. It is best to orient speakers away from surfaces that scatter or diffuse the sound. Avoid placing speakers behind structures that cause distortion, echoes, or sound shadows. When sound localization is used to direct the operator to particular alarm display devices, the sound sources should be oriented so that their location is quickly discerned and corresponds to the location of the intended alarm display device.

The intensity of auditory signals should be such that they are reliably detected under the most adverse, anticipated, background noise conditions. To guard against the possibility that operators inadvertently reduce the audio level so as to render the signals inaudible, the systems that generate auditory signals typically do not allow operators to adjust the signal level. Unfortunately, it is not unusual for signals to be so intense that they irritate or startle the operators. Consequently, operators may turn off or disable alarm systems. Improved approaches to auditory signal design and level selection alleviate this problem.

A signal level 10 dB(A) above average ambient noise is generally considered adequate to ensure that a signal is audible. dB(A) refers to decibels as measured using one of three standard weighting networks (designated A, B, C) typically available in sound-level meters. A-weighted sound levels are typically used when the effects of sound on people are of interest because they correlate well with perceived loudness and speech interference. It has also been recommended that sound intensity should be limited to a maximum of 95 dB(A) but that signal levels of 115 dB(A) may be used if considered absolutely necessary to guarantee effectiveness for alarms indicating extreme danger. Levels this intense are probably appropriate only for situations requiring prompt evacuation of an area. Techniques exist to support the design and production of signals that more effectively convey alarm-related information. Recently developed approaches to auditory alarm signal design recommend that the intensities chosen for auditory signals take into account the frequency spectrum of the signals and of the background noise (16,17). Because of the nature of the human auditory system, signals need not be very intense to be heard reliably as long as the frequencies at which their energy is concentrated differ from those occupied by the background noise. It should be recognized that auditory signals themselves may contribute to the background noise, that is, it may be necessary to consider the audibility of a signal not just in the presence of ambient noise but also in combination with other signals that might plausibly occur at the same time. To avoid mutual masking, the frequency spectra of auditory signals associated with alarms that may be active at the same time should be different. Interference among alarm signals is less of a concern when the signals consist of a number of widely separated frequency components (rather than a pure tone) or of brief groups of pulses presented at intervals (rather than a continuous tone).

It is not good design practice to code auditory signals by intensity (loudness). The range of intensities between the level required to ensure audibility and the level at which signals become aversive is relatively narrow. Therefore, the usefulness of this dimension for coding is limited. If such coding must be used, use only two levels, and distinguishing the signals from each other by a minimum of 6 dB(A). Whether this coding is effective depends on the frequency spectrum of the ambient control room noise and the frequency of the signal.

Each auditory signal should be unambiguous and easily distinguishable from every other tone in the control room. For example, the auditory signal associated with SDCV alarm displays should be easily distinguishable (based on signal characteristics or sound source) from the auditory signal associated with an alarm message displayed by other means (e.g., on a VDU message display). Auditory signals used for new alarms should be separate and distinct from tones used to signify clearing alarms. The latter can be momentary or "self-silencing." If the tone indicating an unacknowledged alarm automatically turns off after an interval of time, a reminder tone can be used to alert the operator to its continued presence.

Current techniques allow designing alarm signals that make better use of the operator's ability to process audio information. It is possible to design signals that are more discriminable from one another than conventional signals and also have the potential to carry more information (18). Sets of readily distinguishable signals can be designed by varying fundamental frequency, harmonic structure, and temporal patterns.

When information is coded by pitch, it is best not to use more than three frequencies. The frequencies should be widely spaced between 500 Hz to 3,000 Hz, although a wider range from 200 Hz to 5,000 Hz may be acceptable. Avoid frequencies in a ratio of 2 : 1 with one another, because it is difficult to identify pitches an octave apart. Signals with multiple frequency components ("chords" or frequency-modulated tones) are more resistant to masking and more easily localizable than pure tones. (If the location of a source sound is to be used as a cue, the signal should not be a high-frequency tone, because such signals are difficult to localize.) No more than three modulated frequency codes for audible alarms should be used. Warbling sounds, with frequencies modulating from one to three times per second, are attention-getting and easily recognized, whereas slower modulation rates do not develop distinguishable characteristics rapidly enough to be appropriate for alarm applications. If modulation of frequency (Hz) of a signal is used to denote information, the center frequencies should be between 500 Hz and 1000 Hz.

No more than three pulse repetition rates should be used for coding purposes. Repetition rates should be between one and eight pulses per second because faster rates may not be perceived as pulses. It has been recommended that repetition rates differ by at least a factor of 2 to ensure operator discrimination. Recent alarm design research suggests that otherwise distinctive signals having similar temporal patterns are readily confused and that this effect may extend to duty cycle. Accordingly, temporal patterns other than regular on/off cycles should be considered.

If sequences of tones are used to represent information, the patterns should be easily recognizable. Warning sounds consisting of "bursts" composed of five or more brief pulses (about 0.1 s in duration) with inter-pulse intervals of 0.15 s to 0.3 s are currently recommended (see Ref. 16). The pulses may be designed to be distinctive with respect to their onset and offset shaping, fundamental frequency, and harmonic structure. The bursts may vary as to the number of pulses, the tempo at which they are presented, and the rhythmic and pitch contours.

It has been recommended that a maximum of nine auditory signals be used when two or more dimensions are used in coding. When signals differ in two or more dimensions (e.g., pitch and temporal pattern), a greater number of signals can be reliably distinguished. This maximum includes auditory signals used outside the control room (e.g., fire alarm or site emergency alarm). The number of conditions for which reliably recognizable audio codes can be used is maximized by taking advantage of differences in the perceived urgency of warning sounds (19) to represent, for example, varying alarm priorities.

Shared Alarms. A single "shared" indicator is sometimes used to represent more than one alarm condition. For example, a "trouble" message may combine several potential problems associated with a single component, or it may address the same problem for a group of similar components (e.g., a bearing temperature alarm may address bearings from more than one component). The types of alarms that may be considered for combination include

- Alarms for the same condition on redundant components, or logic trains, when each has a separate indicator and the indicators are placed in close proximity on the console (e.g., pump *A* or *B* trip, logic train *A* or *B* actuation)
- Alarms for several conditions relating to one component or several redundant components, which require the operator to obtain further diagnostic information either by sending an auxiliary operator out to the component(s) or checking the computer (e.g., pump *A* or *B* trouble)
- Alarms for several conditions that call for the same corrective action
- Alarms that summarize single-input alarms elsewhere in the control room

Because they require additional effort by the operator to identify the specific alarm, it is best to minimize the number of shared alarms, especially when

- Different actions are to be taken depending on which alarm condition exists and the operator cannot readily determine which constituent is alarming
- The required response must be initiated immediately, so that taking time to determine which constituent is alarming would risk an inadequate response
- The operator's understanding is improved by indicating the conditions separately because of similarity to the layout of associated controls
- The constituent conditions are not of a similar nature or are not of the same order of importance, so that the action to be taken is very different depending on which condition is alarming

If a new parametric deviation has occurred before a preceding alarm has cleared, the shared alarm should return to the new alarm state (e.g., reflashing). It is also important that the alarm system enable the operator to reactivate the visual and audible alert indications for the alarm when subsequent alarm conditions occur after the initial alarm condition has been acknowledged.

Organization of Alarms. Grouping alarms within a display by function, system, or other logical organization facilitates the operator's understanding of the relationships among alarms and among alarms and system components. Functional groups should be visually distinct from one another. Clear labels and delineations among groups allow the operators to determine easily which systems have alarms. If alarm displays are organized in matrices, the vertical and horizontal axes of the displays can be labeled with alphanumerics so that a coordinate designation is available for any particular visual element. Coordinate designation is preferred on the left and top sides of the display. The size of an alarm tile matrix should be limited to 50 alarms.

Operators use alarm displays more effectively if alarms are arranged according to naturally occurring relationships, such as those derived from the physical process, for example,

- Alarms for a given thermodynamic parameter at different points within the system which indicate a progression (e.g., within a fluid system, a series of pressure alarms starting with the source tank and ending with the system discharge) could be arranged left to right
- Several alarms for the same variable indicating levels of severity (e.g., tank level low and tank level low-low) could be arranged in a vertical array
- Alarms related by cause and effect could be adjacent to one another

Once an arrangement has been chosen, the arrangement can be used consistently within similar systems or alarm groups.

Alarm message lists are more effective when segregated by alarm priority with highest priority alarms listed first. In addition to priority grouping, provide operators with the capability to regroup alarm messages according to operationally relevant categories, such as function, chronological order, and status (unacknowledged, acknowledged/active, cleared). For example, it can be useful for diagnosis to arrange alarms in chronological order with the most recent messages placed at the top of the stack. A separation (blank row) every four or five alphanumeric messages enhances readability.

Alarm Controls

Alarm systems typically include controls to: (1) *silence* the auditory component of the alarm, (2) *acknowledge* the meaning of the alarm, (3) *reset* the alarm to its monitoring state, and (4) *test* the alarm display characteristics. Making these controls easily distinguishable from each other by touch and sight helps prevent accidental operation of the wrong control. Techniques, such as color- or shape-coding of individual controls and color shading or other demarcation of groups of alarm controls, can be used.

Silence. Operators sometimes disable distracting or irritating auditory signals. This can be a dangerous situation. Good auditory signal design, such as described in the previous section, helps minimize the need for such action. Another solution is to provide the capability to silence an auditory signal from any set of alarm system controls in the primary operating area. Although manual silence is a generally desirable feature (in that it increases the likelihood that the operator has attended to the alarm information), it may become distracting to silence all alarms manually under high-alarm conditions.

Alarm system designs should not allow the operator to defeat the control. For example, some pushbuttons used for alarm silencing and acknowledgment can be held down by inserting an object in the ring around the pushbutton.

Acknowledge. Acknowledgment terminates the alarm flashing and is usually indicated by steady illumination until the alarm is cleared. Acknowledgment should be possible only from locations where the alarm message can be read. If alarm information is available at multiple VDUs, then operators should be able to acknowledge the alarm from the VDU at which they are working.

Reset. The reset control places the alarm in an inactive state after an alarm condition has cleared. When it is important to inform operators explicitly of a cleared condition, a manual reset is appropriate. An automatic reset option is useful when operators have to respond to numerous alarms or when it is essential to reset the system quickly. The reset control should be effective only from locations at which plant personnel know which alarm they are resetting.

Test. Given its importance, it is desirable for the alarm system to indicate positively to the operator when alarm system malfunctions occur. By the same token, periodic testing of the system by operators is good operational practice. Test controls should be available to initiate operational test conditions for all essential aspects of the alarm system, including processing logic, audible alarms, and visual alarm indications (e.g., to detect burnt-out lamps).

Beyond these silence, acknowledge, reset, and test controls, computer-based alarm systems may require other controls to allow operators to sort the alarms according to time or component, to define temporary alarms, to adjust set points, and to control processing options.

Automated and Modifiable Characteristics

In certain situations, such as major system disturbances, it may be desirable to reduce operator workload by automating or modifying some alarm system functions. For example, lower priority alarms might be automatically silenced, or the flashing of unacknowledged alarms temporarily stopped. Similarly, automated controls may be implemented to trigger appropriate displays, such as alarm graphics, data windows, or display pages. Other dynamic aspects of the alarm system may allow operators to introduce operator-defined characteristics, such as alarm parameters and set points.

If the alarm system automatically changes operational configurations under some situations, an alert is needed to indicate that the configuration has changed. Alarm systems may provide the capability for operators to select alternative functional configurations under some alarm situations, such as automatic silence of auditory alerts for lower priority alarms under high-alarm conditions. It is important for the alarm system to indicate to operators that a requested change in system configuration has been successfully achieved. In addition, a prominent display of the present configuration should be available.

Requiring operator confirmation of any significant change in the alarm system, whether selected by the operator or automatically, prevents potential errors resulting from an operator's misreading of the alarm system's present configuration.

The alarm system may provide temporary, operator-defined alarms and operator-defined set points for specific conditions (e.g., temporary alarms to support increased monitoring of a problem component or of a parameter trend that is approaching a limit). A clear indication of operator-defined alarms and set points is needed that is distinct from the alarm/set points designed into the system. It is also important that operator-defined alarms and set points not override or interfere with the existing alarms and set points.

Reliability and Maintenance

It is important that the hardware and software components of the alarm system are sufficiently reliable to prevent single component failures from causing significant loss of functions or information. For example, redundancy and diversity of the alarm system design can protect against losing alarm indications or generating spurious alarm messages as the result of sensor or signal processing malfunctions.

Tile-type displays can be designed with dual light bulbs so that a single bulb failure does not interfere with the operator's detection of the alarm condition. In case of flasher failure of an active alarm element, the element should assume a highly salient state, such as a high flash rate or a steady on (e.g., illuminated) state rather than a less salient state, such as off. Although it is preferable in a flasher failure that the active alarm element remains on (e.g., illuminated) rather than off, a unique and highly salient code is best. In addition, other alerting mechanisms, such as warning messages, may be used to inform the operator of a malfunction in the alarm display system.

Where VDUs are the primary means of displaying alarms, additional reliability is gained by making it possible to access the alarms from more than one VDU. Alarm printers also back up VDU displays.

The alarm system should be designed so that maintenance can be performed with minimal interference with the activities of the operators. Desirable features include built-in test capabilities, modular components that can be rapidly removed and replaced, and rear access panels which prevent maintenance activities from obstructing the operator's view of controls and displays.

When an alarm is taken out of service for maintenance, its visual and audio signals should be disabled and cues provided to indicate clearly that the alarm is out of service.

Location and Integration

Locate visible alarm elements within about 60° on either side of the direct line of sight of the operator's normal work position. To avoid confusion, do not locate alarms near indicator lights that present information about the state of equipment. Alarm displays and controls should be arranged and located so that the operators who must respond to an alarm can access the alarm information quickly enough to respond adequately. Avoid arrangements that require one operator to read an alarm message only to recite it to another person. It is also important to position alarms near the other displays and controls that are required for diagnostic or corrective action in response to them.

Consistency between the alarm system and the other displays and controls in the workplace is important to minimize human error. Thus, the same design conventions for elements, such as symbols, icons, acronyms, coding, and measurement units that are used in other operator interfaces and procedures, should also be applied to the alarm system. For example, if color is used to code priority, the relationship between the colors used and level of priority should be the same in the alarm system as in the process displays. For example, if red is used to indicate the highest priority alarm, then red should also be used to indicate high-priority information in the process displays.

Alarm Response Procedures

Alarm response procedures (ARPs) provide more detailed information concerning the nature of the alarm condition than is typically provided in an alarm message. They are especially important to operators when an unfamiliar alarm is activated or when an alarm seems inconsistent with the operator's understanding of the plant state.

Operators should have immediate access to ARPs from the location at which the alarm messages are read. ARPs may be hard copy or computer-based. In a tile-based alarm display system, the operator's access to ARPs can be aided by identifying and indexing the ARPs consistent with the method of identifying the alarm, for example, by row and column designations.

A good ARP contains the following information:

- The system/functional group to which the alarm belongs
- The exact alarm text or legend
- The alarm source (i.e., the sensor(s) sending the signal, including processing or signal validation and conditioning logic, and the actuating device(s) for the alarm with a reference to a schematic diagram on which such devices can be found)
- Alarm set points
- Priority
- Potential underlying causes for the alarm (e.g., low water level or feed flow deficient in the long term)
- Required immediate operator actions, including actions the operator can take to confirm the existence of the alarm condition
- Actions which occur automatically when the alarm occurs (and which the operator should verify as having taken place)
- Follow-up actions
- Pertinent references

Just as alarm design conventions should be consistent with those applied to other interfaces, the information in the ARPs should reflect the same conventions applied elsewhere.

ALARM SYSTEM EVALUATION

This section considers the general evaluation of the alarm system with regard to its functional objectives. It does not address engineering tests of such items as the correctness of wiring or the performance of the circuitry.

The objective of the alarm system review is to ensure that the alarm system functionality and its associated controls, displays, and data processing support safe, efficient, and reliable operator performance. The following design evaluations can be performed to help ensure that the alarm system meets its objectives:

Task Support Verification. This evaluation verifies that all necessary system alarms are provided in the alarm system. The method involves comparing a list of required alarms to the alarms available in the workplace. A problem is identified if (1) required alarms are not provided by the alarm system or (2) alarms are present in the workplace that have not been identified as required.

Human Factors Engineering Design Verification. This evaluation verifies that the alarm system design and implementation take operator capabilities and limitations into account. The method involves evaluating alarm design characteristics against human factors engineering (HFE) standards and guidelines for alarm systems, such as Ref. 20. Problems are identified if the design is inconsistent with HFE guidelines.

Integrated System Validation. This evaluation validates that the integrated alarm system design supports operator task performance. This type of evaluation is best performed using an actual system or prototype under realistic operational conditions. When this is not practical, real-time, dynamic simulation of a system may provide an acceptable test bed. Dynamic performance evaluation addresses both (1) the operator interfaces associated with operation of the alarm system and (2) the quality, accuracy, timing, and usefulness of the information provided by the alarm system to plant personnel. Problems are identified if task performance criteria are not met or if the alarm system imposes a high workload on plant personnel.

Then problems identified through these evaluation activities can be remedied prior to actual operational use, resulting in an effective alarm system that helps operators to monitor the system and to detect disturbances in a timely manner.

BIBLIOGRAPHY

1. D. Green and J. Swets, *Signal Detection Theory and Psychophysics,* New York: Wiley, 1988.
2. R. Sorkin and D. Woods, Systems with human monitors: A signal detection analysis, *Human Comput. Interaction,* **1**: 49–75, 1985.
3. R. Sorkin, B. Kantowitz, and S. Kantowitz, Likelihood alarm displays, *Human Factors,* **30**: 445–459, 1988.
4. J. G. Kemeny, *Report of the President's Commission on the Accident at Three Mile Island,* Washington, DC: US Government Printing Office, 1979.
5. J. Rogovin, Three Mile Island, *A Report to the Commissioners and to the Public,* Washington, DC: US Nuclear Regulatory Commission, 1980.
6. E. Marshall and F. Owre, The experimental evaluation of an advanced alarm system, in *Advances in Human Factors in Nuclear Power Systems,* La Grange Park, IL: American Nuclear Society, 1986.
7. Y. Fujita and T. Sanquist, Operator cognitive processes under abnormal plant conditions with conventional and advanced control room designs, *1988 IEEE 4th Conf. Human Factors,* New York, 1988.
8. Y. Fujita, Improved annunciator system for Japanese pressurized-water reactors, *Nuclear Safety,* **30**: 209–221, 1989.
9. A. Ephrath and L. Young, Monitoring vs. man-in-the-loop detection of aircraft control failures, in J. Rasmussen, and W. Rouse, (eds.), *Human Detection and Diagnosis of System Failures,* New York: Plenum, 1981.
10. N. Moray, The role of attention in the detection of errors and the diagnosis of failures in man-machine systems, in J. Rasmussen and W. B. Rouse (eds.), *Human Detection and Diagnosis of System Failures,* New York: Plenum, 1981.
11. R. Sorkin, Why are people turning off our alarms?, *Human Factors Soc. Bull.,* **32** (4): 3–4, 1989.
12. J. Reason, Generic error-modelling systems (GEMS): A cognitive framework for locating common human error forms, in J. Rasmussen, K. Duncan, and J. Leplat (eds.), *New Technology and Human Error,* New York: Wiley, 1987.
13. J. Reason, Modelling the basic error tendencies of human operators, *Reliability Eng. and Syst. Safety,* **22**: 137–153, 1988.
14. D. Kahneman and A. Triesman, Changing views of attention and automaticity, in R. Parasuraman and R. Davies (eds.), *Varieties of Attention,* New York: Academic Press, 1984.
15. K. Bennett and J. Flach, Graphical displays: Implications for divided attention, focused attention, and problem solving, *Human Factors,* **34**: 513–533, 1992.
16. R. D. Patterson, *Guidelines for Auditory Warning Sytems on Civil Aircraft,* CAA 82017, London: Civil Aviation Authority, 1982.
17. C. LaRoche et al., 'Detectsound': A computerized model for predicting the detectability of warning signals in noisy environments, *Appl. Acoust.,* **33**: 193–214, 1991.
18. J. Edworthy, S. Loxley, and I. Dennis, Improving auditory warning design: Relationship between warning sound parameters and perceived urgency, *Human Factors,* **33**: 205–231, 1991.
19. J. Edworthy, Urgency mapping in auditory warning signals, in N. Stanton (ed.), *Human Factors in Alarm Design,* Bristol, PA: Taylor & Francis, 1994, pp. 14–30.
20. J. O'Hara et al., Human Factors Engineering Guidelines for the Review of Advanced Alarm Systems, NUREG/CR-6105, Washington, DC: Nuclear Regulatory Commission, 1994.

Reading List

Annunciator Sequences and Specifications (ANSI/ISA-1979). Research Triangle Park, NC: Instrumentation Society of America, 1992.

J. Edworthy and A. Adams, *Warning Design: A Research Prospective,* London: Taylor & Francis, 1996.

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

MPR Associates, *Power Plant Alarm Systems: A Survey and Recommended Approach for Evaluating Improvements,* EPRI NP-4361, Palo Alto, CA: Electric Power Research Institute, 1985.

J. O'Hara et al., *Human factors engineering guidelines for the review of advanced alarm systems* (NUREG/CR-6105), Washington, DC: US Nuclear Regulatory Commission, 1994.

J. O'Hara et al., *Human-system interface design review guideline* (NUREG-0700, Rev. 1), Washington, DC: US Nuclear Regulatory Commission, 1996.

R. D. Patterson, *Guidelines for Auditory Warning Systems on Civil Aircraft,* CAA 82017, London: Civil Aviation Authority, 1982.

N. Stanton, *Human Factors in Alarm Design,* London: Taylor and Francis, 1994.

JOHN M. O'HARA
WILLIAM S. BROWN
Brookhaven National Laboratory