ABNORMAL SITUATION MANAGEMENT:
NOT BY NEW TECHNOLOGY ALONE...

Edward Cochran
Peter Bullemer

Honeywell Technology Center
3660 Technology Drive
Minneapolis, MN 55418

Abstract

Abnormal Situations have always challenged operations personnel, and they likely always will. Abnormal Situation Management is a particular challenge at this point in history because the aggressive application of increasingly complex processes, sophisticated control strategies, and highly integrated approaches to production planning have led to productivity levels only dreamed of by previous generations of process engineers. However, the capacity of human operators to deal with this complexity, and the sophistication of their tools, has remained essentially unchanged: Operators still rely on 10,000-year old perceptual and cognitive capacities and user interface ideas from the 1980s. This situation will change greatly in the next few years: **There are a number of emerging technologies that address key aspects of the human side of the ASM problem, and all of them must be used in a comprehensive solution. The remaining technology development and transfer challenge will be to ensure that these user support technologies can be introduced into the current operations culture in a way that results in it, too, evolving to the next paradigm for success**.

The ASM Problem

Abnormal Situations comprise a range of minor to major process disruptions in which operations personnel have to intervene to correct problems with which the control systems can not cope. Preventable losses from Abnormal Situations cost the U.S. economy at least $20B annually—about half of that in direct losses to the petrochemical industries. Increased demands for higher efficiency and productivity in these industries are resulting in tremendous increases in the sophistication of process control systems through the development of advanced sensor and control technologies. However, these sensor and control technologies have not eliminated abnormal situations and will not in the future. Consequently, operations personnel continue to intervene to correct deviant process conditions.

The persistent paradox in the domain of supervisory control is that as automation technology increases in complexity and sophistication, operations professionals are faced with increasingly complex decisions in managing abnormal situations. A contributing factor to this phenomenon is that the sophistication of the user support technologies has not kept pace with the task demands imposed by abnormal situations. Thus,

Paper presented at the AICHE 1996 Safety Conference   Page 1

collaborative decision support technologies must be developed to significantly improve abnormal situation management practices.

Designing ASM Systems to Solve the Problem

A comprehensive approach to the design of the human-machine system interaction is needed so that operations personnel receive information appropriate to their needs, while at the same time appropriate members of the operations staff are able to collaborate to solve the problem as a team. Individual needs vary as a function of a large number of variables: the current situation, the task being performed, individual preferences and styles—and others yet to be determined. In order to serve these needs, we need to carefully assess the information requirements, not just for the current job functions present in existing plants, but for the job functions that will evolve as better decision aids become available and operators receive more support.

In the aviation community, the introduction of information technology and decision support tools resulted in nothing less than a redefinition of the jobs of the flight crew: Whereas good hand-eye coordination, excellent spatial cognition skills, and well-honed aviation intuitions used to define a large part of the repertoire of talented pilots, the current job benefits more from a thorough understanding of automation systems, good procedural discipline, and cockpit resource management skills.

If we successful in applying user support technologies appropriately to enable effective abnormal situation management, the role of the process operations staff will change just as significantly. The trick will be to introduce these technologies in a way that is simultaneously consistent with the existing culture and enabling of the new one.

The key to success is, to refrain from identifying applications for the latest technological advances. Instead, we need to identify the problems that have to be solved and only then search for answers in both, sociocultural and technological terms. In the remainder of this paper, we'll describe several examples of how this approach can be used. We'll start with challenging but solvable problems—problems for which an approach is reasonably well-defined and eventual success is reasonably likely. We'll end with the extremely difficult problems, for which consensus does not exist even for an initial approach.

Examples of problems, unmet needs, and solution concepts

Problem: Operations and maintenance personnel spend too much time and effort seeking information

When the process is running well, as much as half the time of a board operator may be spent communicating with field personnel[1], typically via two-way radio. The communications task imposes a tremendous drain on the attention of all individuals involved, yet the amount of information that can be transferred is relatively small, the initial reliability of the transmission is often uncertain (resulting in needless repetition) , and the information is rarely presented in the most appropriate form for the user.

Operations and maintenance personnel in the field need more information, more rapidly, more reliably, and with less involvement of control room personnel. Information could include equipment configurations, current status, calibration and maintenance histories, schedules, work orders, and other operations data. Control room personnel[2] need more information about the field, as well: locations of people and other resources, status of ongoing activities, and confirmation/investigation of readings in the control room.

---

[1] Field personnel include job classes such as field operator, chief operator, mechanic, and instrument technician.

[2] Control room personnel include job classes such as board operator, shift supervisor, maintenance coordinator, process engineer and operations engineer. Some personnel work in both the control room and field depending on the task such as chief operator, shift supervisor or maintenance coordinator.

Technologies are becoming available that can be applied to this problem. Instead of garbled walkie-talkie voice communications, high-bandwidth RF networks can provide reliable, secure, authenticated digital communication between field operators and the control room. Options include commercial cellular, commercial digital packet, spread spectrum, direct satellite, infrared (with wire/fiber trunks), and bi-directional paging systems.

Gateways to process data can be linked to these networks, enabling field access to process data.

Maintenance information can be converted to digital databases, tagged with a standard indexing scheme such as SGML, and provided to the field on an as-needed basis. "As needed" can mean for the right operator, at the right time, and in a way consistent with their knowledge, preferences, and operational authority.

The information can be presented on portable displays, hand-held "digital assistants", clipboard-sized portable maintenance terminals, wearable computers, and visor- or head-mounted displays.

Given such systems, new sensor technologies such as uncooled solid-state IR imagers, UV sensors, and photo-multipliers can be used to create augmented vision systems to provide information heretofore unavailable. The information in CAD databases can be combined with operational data and sensor information to create virtual displays. The location of field operators can be made available in real time to dispatchers using GPS systems or information from the RF network itself.

Given advanced video and remote sensing capabilities in the field, the console operator could take advantage of them, as well: They could view the same scene as remote operators, and be able to coordinate operations in real time. They could confirm some sensed values remotely with real time imagery (e.g., tank levels), saving time in abnormal situations.

Providing all of these capabilities all at once is just becoming technically feasible, but it is not yet cost effective—and may never be. And without a significant expected benefit, the changes in the operating cultures these technologies will require—and engender—may represent an obstacle to their adoption. If history is any guide, we could see field use begin in a particular high value area, and expand rapidly as the operations culture changes as operators become more effectively supported. To enable this evolution to occur as rapidly as possible, we should strive for modularization of the technology components, incremental development of the most comprehensive solutions, and focused effort on training to support the reengineering of the operator's role.

Problem: Operations personnel think about problems, DCSs deal with points and data values.

An operator faced with a known problem—a compressor trip, a malfunctioning instrument, a failed pump—typically manages the situation well. A more difficult situation is when the operator knows a problem is occurring, but has not yet diagnosed it. Most difficult of all is the situation in which the operator is not even aware that a problem is occurring.

From the operator's perspective, the definition of "problem" depends on the context: A low flow situation can be an emergency condition, or it can be an expected condition during a process start-up or shutdown. Current DCSs typically are not configured to address any of these distinctions.

Instead, operations personnel are presented with process data in the form of PVs, and problem symptoms in the form of alarms based upon preset limits with respect to those PVs, and the rest is up to them. Alarms are often not context sensitive, so systems that are finely tuned to support optimization may generate thousands of alarms during a process upset. Alarms are rarely associated with specific problems on a one-to-one basis, and DCSs may be entirely mute about some situations that need attention (e.g., a stuck sensor, or miss-communicated lock-out status).

DCSs are inherently compatible with problem-based alarming: Mass balance analyses, expert systems, and statistical diagnostic techniques are becoming more widespread, albeit very slowly. What's needed to accelerate this trend is better ways to combine and aggregate data, better tools for easier, perhaps even

automatic, development of such problem monitors, and higher-level, more comprehensive representations for plant equipment and processes.

General solutions are becoming available to fill these needs. Software infrastructures are being developed which will permit multiple applications to diagnose problems based upon shared information and cooperate to plan the best course of action. Agents, authoring environments and tools, comprehensive plant data reference models, and open operating environments will help solidify these gains.

Nevertheless, as with the previously discussed problem example, the fact that a technical solution may be feasible—even available—will not guarantee its widespread adoption. Careful attention to the concomitant impact of the technology, and the requirements to minimize that impact—is required. The operations culture, alarming methodologies, process safety management practices, and software maintenance infrastructure, among others, will be greatly affected by the introduction of these technologies, and those who ignore this fact will delay their adoption, regardless of the economic benefits.

Problem: The Operator/Machine System interaction model is not well-defined, let alone standardized

Fifty years of human factors studies in the aviation industry have led to the adoption of automation systems whose design focus is squarely aimed at supporting pilot-machine system interactions that have been carefully developed in a principled and rigorous way.

Twelve years of product evolution in the desktop computer industry have led to systems in which user-machine system interaction metaphors vary wildly, conflict both within and across applications, and often fail to support the user's expectations in fundamental ways.

The DCS represents a hybrid of these outcomes: In most cases, DCS vendors anticipated and supported the expectations of their users, even though these users had been working with pre-digital technologies. And, early and careful attention to human factors issues led to the reasonable support of DCS-specific tasks as performed by these same users. Problems arose as three trends simultaneously developed:

- DCSs evolved into much more complex entities, networked with a variety of outside systems, supporting more complex control strategies, and so on, and

- Processes grew more complex as producers worked to increase efficiency and reduce costs, and

- Everyone became observers of, if not participants in, the desktop computing revolution, which has set expectations (good and bad) for what interaction with a computer should be like.

The result is that a pastiche of methods has been created, and the users have to adapt to them as they change tasks, systems, roles, and so on. The situation is comparable to the aviation industry in the early years, in which one aircraft's lever for raising flaps was another's for lowering landing gear—with predictable results.

This issue has to be addressed, and continually readdressed; we must not settle for today's state of the art tomorrow.

A comprehensive model for user-machine system interaction in the process industries must be rigorously developed, tested, challenged, and revised, and a standard of practice must be established to support it. Systems must evolve so that the operator is not routinely swamped with information, aggravated by the user interface, required to use error-prone techniques to enter data, or exposed to situations in which being misled is even a remote possibility. The system must completely prevent adverse consequences from happening when the interaction of individuals predictably leads to misunderstandings, misperceptions, and mistakes. It must also reduce, by orders of magnitude, the level of what post-incident review teams always label "human error."

There should be no such thing as a break down in lock-out, tag-out procedures—the user-machine system interaction model should utterly prevent such things from being possible.

There should never again be colored text on clashing colored backgrounds on operational displays—the user interface development tools should make it very clear to the developer why such a design is inappropriate.

Users should never again have difficulty in navigating from one display to another, be able to enter a set point outside the controller's capabilities, or perceive the data from one unit as coming from another.

And, looking to the future, decision support systems must never act like a back-seat driver when what the user needs is a helpful child—or vice versa.

Solution components for this problem are also beginning to emerge, but there is little consensus yet as to how to apply them. Operator intent recognition can help systems act in task-specific ways. Task modeling can help online information systems provide relevant (as opposed to canned) support. Tailored user interface displays can ensure that color-deficient users can differentiate key data, users preferring graphs can see lots of graphs, and users needing quantitative can see lots of appropriate numbers. And, user-centered design methodologies can ensure that we address this whole problem area in an empirically rigorous way in when the analytically rigorous methods are lacking.

Problem: The human behavior that can lead to abnormal situations is not changed merely by exposure to information; the human culture that is the focus of technology transfer efforts is not changed merely by exposure to user support technologies designed to improve ASM practices.

This is the most challenging problem of all. Much of the solution to the ASM problem requires no new technology—"merely" the rigorous application of solutions that have already been developed. But the ASM problem remains, because changing the standards of practice in a human endeavor is far more difficult than introducing the new technical solutions.

For example, the U.S. Navy, over the past decade, has reduced flight-deck incidents by x% through a comprehensive effort to instill absolute procedural discipline. Little technology was involved, yet the effort still required a decade.

The availability of the most modern medical prenatal technology available has had little effect in many areas of the U.S., because the potential users of the technology have not had access to, or perhaps do not avail themselves of, that technology.

Recognition that the ASM problem costs the industry at least $10B annually in the U.S. alone is the first step. Commitment to change—to no longer accepting that cost as part of doing business—is the next. Acknowledging that there are no silver bullets—no easy cures—is the third. User support technologies necessary to solve the problem are becoming available; integration of solution components into the existing systems is underway.

The operations culture will need to change more rapidly than it has in the past, and the industry will have to evolve accordingly. Procedure systems, configuration management systems, standards of practice, sharing of experiences, automation of procedures, effective training systems, and management support will all help, but, as with everything else, the most effective change agents will be the people who realize that there is a better way.

Conclusions

The adoption of new user support technologies will require a change in the process industry's culture; need to be careful not to plan/design/buy these systems entirely with respect to current models of doing business.

By some measures of thinking, we are currently enjoying the third generation of desktop computing systems. This technology evolution has taken place within a mere fifteen year span. The process industry tends to change that less rapidly, primarily because of the capital expense and system complexity involved.

Indeed, the process industry is by some measures also enjoying the third generation of process control technology, but it has taken us about 75 years to get here.

If new user support technologies for ASM practices can be as successful as their aviation counterparts, the fourth generation of control systems will be enabled. Since no one can predict the changes that will result [recall that the third generation control system—the DCS—gave us the ASM problem in the first place], it is imperative that we develop our approaches to be as flexible, modular, and incrementally upgradeable as possible.

Since the culture is required to change, focusing the technology development and transfer approach on addressing the capabilities of the future users is critical.

The great constant in the evolution of complex human-machine systems remains the capabilities and limitations of the human users. Designing new systems with those capabilities and limitations in mind is absolutely necessary—albeit not sufficient—for the success of those systems. The introduction of technology that can not be used is pointless, regardless of the new capabilities thus introduced. Once in a while, technology developers rely on their usability intuitions in designing their systems, and succeed—but history is replete with examples of less positive outcomes. If the system can be used, and provides benefit, the user can and will adopt it, changing their culture in the process. This is the outcome we desire in improving abnormal situation management practices, even if it does make the technology development and transfer task a bit more challenging.