Lintern, G. (1992). Risk management: The irremedial incompleteness of rules. *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics* (pp. 833-836). Piscataway, NJ: IEEE.

RISK MANAGEMENT: THE IRREMEDIAL INCOMPLETENESS OF RULES

Gavan Lintern
Aviation Research Laboratory and Beckman Institute
University of Illinois at Urbana-Champaign
Savoy, IL 61874

Abstract **Risk management poses a major challenge in the design and operation of large-scale industrial systems. The common strategy for managing risk is the rule-based approach in which automatic safety devices and standard operating procedures are developed as a counter to the potentially catastrophic effects of system malfunctions that can be anticipated. This is, however, a flawed approach. Large-scale industrial systems can be so complex that it is impossible to anticipate all eventualities. In addition, rules are irremedially incomplete as determinants of behavior. They encapsulate explicit knowledge but they cannot be used to completely specify behavior because their use always involves implicit or unarticulated knowledge. A case study of a highly reliable system with catastrophic potential is described. The development of experiential knowledge is seen as a key element in the achievement of safe and reliable procedures. The case study suggests that operators should be given direct access to critical information about the process to be controlled and should be provided with ample opportunity to develop sensitivity to this critical information.**

INTRODUCTION

Large-scale industrial systems such as chemical factories, power generation plants, and transportation have become essential components of our social order. These systems do, however, impose an element of risk in their potential for catastrophic breakdown; the nuclear power explosion at Chernobyl, the loss of oil in Prince William Sound from the Exxon Valdeez, and the release of poisonous gas from the Union Carbide chemical plant at Bhopal providing a small sample of the more heavily publicized events. The impact on human health and the environment of breakdowns in these systems can rival the impact of natural disasters such as drought, fire, and flood. In some cases the consequences can linger for decades after the event.

From one perspective, breakdowns in large-scale, industrial systems are inevitable. It is unrealistic to expect that we can insulate such systems from malfunctions. Thus, it might be argued that risks are unavoidable, and that the proper approach is to assess the risk/benefit tradeoff. There is, however, an emerging view that the risks are due less to the inherent potential of the process to inflict damage than to inappropriate design and management practices [1, 2]. Others have argued that a pervasive cultural emphasis on the primacy of explicit knowledge and rational thought has driven design and management towards more elegant and efficient but more brittle and error-prone systems [3, 4, 5]. The risks inherent in the operation of large-scale industrial systems may be accentuated by limited design philosophies [6] and by poor management practices [7].

THE FALLACY OF OPERATOR ERROR

Many industrial disasters have involved actions by operators at the human-machine interface. In most cases it could be said that the disaster would have been averted had the operators acted differently. In retrospect, it is clear that relatively ordinary control actions by operators could have prevented the nuclear power incident at Metropolitan Edison's Three Mile Island plant in 1979. These operators were subsequently criticized for their performance and one member of the Kemeny commission appointed to investigate the incident accused them of stupidity [1]. Similarly, actions by operators were involved in the 1984 chemical accident at Union Carbide's Bhopal plant and in the capsize of the ferry "Herald of Free Enterprise" at Zeebrugge in 1987 [2].

Although convenient, a finding of operator error as a cause for a major industrial incident is generally inappropriate. It could only be a valid conclusion if operators were deficient in some aberrant sense that could not be foreseen by those responsible for hiring and training. It is unlikely that any group of operators would have taken the required action within the context of misinformation, obscure interactions, and inadequate control interface within Metropolitan Edison's plant at Three Mile Island [1]. Operator actions that contributed to the ferry accident at Zeebrugge and the chemical accident at Bhopal were, in large part, induced by attitudes and procedures well known to and reinforced by management [2]. Fault in such cases lies not with operators but with systemic problems in design and management. Rather than blaming operators after the fact, the more desirable approach is to implement procedures and to design systems to minimize risks in the event of a failure.

RULE-BASED CONTROL

A common approach to risk management is to develop sets of rules that operators can employ to cope with all anticipated problems. Alternatively, rules instantiated as automatic safety devices may bypass the supposed vagaries of operator behavior by transferring the responsibility of rule implementation to the design stage. To some, rules promise the removal of uncertainty and therefore of risk. Nevertheless, complex rule sets either in the form of instructions to operators or in the form of automatic safety devices have so far failed to insulate us from major accidents. The common response to procedural failures has been to extend the rule set and to emphasize enforcement. For example, the Kemeny commission proposed that nuclear power generation be administered by a para-

military organization that could enforce rules more stringently [1]. There are, however, reasons to believe that continued elaboration of rule sets, or more extensive deployment of automatic safety devices will not have the desired effect.

PROBLEMS OF NOVELTY AND COMPLEXITY

One assumption central to the rule-based approach is that sources of risk can be identified in advance. This has, however, not yet been possible for large-scale industrial systems. Clearly, it is prudent to devote some effort to anticipating sources of risk. Some dangers are obvious and preventative measures straightforward. Automobile safety restraints, domestic smoke detectors, and bicycle safety helmets are simple responses to obvious risks. Other dangers, such as those posed to young children by residential swimming pools [8] may not be so obvious, but again the solution is relatively straightforward.[1] In large-scale industrial systems the potential problems become more ambiguous and more problematic to identify. It is often difficult to discriminate real from illusory dangers in advance and, in fact, the precise circumstances and the course of a serious incident have often never been imagined prior to the event.

In addition, solutions proposed for anticipated dangers can sometimes have a negative impact on risk management. At Three Mile Island, emphasis by management on the dangers of a high pressure incident that might fracture the containment integrity delayed the correct interpretation of this event as a loss of coolant accident. At the Fermi nuclear plant in Monroe, Michigan a partial fuel melt in 1966 was caused by a metal vane that dislodged to block coolant flow. On the insistence of a committee of the Nuclear Regulatory Commission, this metal vane had been installed to guard against anticipated problems resulting from coolant surges. In both of these cases, attempts to reduce risk actually contributed to the problem. More generally, rule-based approaches to risk management are limited in large part because of the considerable potential for complex, large-scale systems to generate novel and obscure interactions. The development of more extensive rule sets and the addition of automatic safety devices only increase the complexity and thereby increase the potential for the generation of even more complex and more obscure interactions.

IRREMEDIAL INCOMPLETENESS OF RULES

A further problem with rule-based control is that rules must be interpreted within the context of situated action. Suchman [5] argues that rules cannot fully specify required control behavior. For any particular situation, a considerable amount of implicit knowledge is required to transform a rule into action. In her terms, there is an "irremediable incompleteness" of instructions. Any actor who is following instructions is inevitably faced with the problems of interpreting those instructions and of ascertaining their practical significance for situated action. There is always room for differing interpretations. Suchman [5] concludes that a pre-established plan as encapsulated in sets of instructions or procedures can act as a resource for action but cannot fully specify it.

In concert with Suchman [5], Winograd and Flores [9] argue that the understanding of language is grounded in unrecognized preunderstandings. It is possible to articulate a preunderstanding to some level, but pursuit of that goal is aided most effectively by "breakdown"; that is an incident in which the unexpected draws attention to an action or interaction that has previously gone unnoticed. For a large-scale industrial system with catastrophic potential, breakdown is an unacceptably risky means of working through potential problems. Furthermore, Winograd and Flores [9] argue that articulation of unrecognized preunderstanding is neverending. The attempt to articulate a preunderstanding is made in terms of language and experiences that themselves reflect unarticulated preunderstandings.

From this perspective, rules can never completely specify an action. There will always be a need for controllors to supplement or interpret a rule set in terms of their own understanding (the irremediable incompleteness problem). To the extent that individuals must interpret a rule within their own body of unarticulated knowledge [9], there will inevitably be local deviations from rule-based precision. It is characteristic of complex, large-scale systems that even minor deviations from ideal behavior can cascade and amplify through the process to generate major or catastrophic breakdowns.

SELF ORGANIZATION ON THE FLIGHT DECK

That complex, large-scale systems can be managed successfully is evident in an analysis by Rochlin, La Porte, and Roberts [10] of flight-deck operations on board an aircraft carrier. A key feature in the development of safe flight-deck procedures is an extended work-up period to prepare for full-capacity operations. After commissioning or refit, weeks are spent in which aircraft launch and recovery rates are gradually increased to operational levels. It seems significant that there are no detailed operating procedures: as noted by Rochlin et al., the only complete operating manual is the working carrier itself. The workup period is critical to the evolution of workable procedures. It is during this period that the crew members develop a functional organization for moving aircraft around the deck and for launching and recovering them.

A unique organization emerges for each carrier, but there is a similitude or self-similar identity at the level of function that is invariant across carriers. That is, each unique organization found in the different U.S. Navy Carriers accomplishes the same job. The role of

[1] The implication of the data presented by Rodgers [8] is that families with young children should not have swimming pools. This may be a difficult choice emotionally but it is not a technologically complex one.

management in this system is to establish global priorities and schedules, but those who execute the multitude of various tasks to accomplish the goals set by management are left to work out their own adaptations to local conditions. An information field is generated and sustained by the action of the workers reacting individually to information in their immediate environment.

The description by Rochlin et al. [10] of carrier operations contrasts with the usual expectations of risk management. To the casual observer, the procedures that emerge will appear to be ad hoc and inefficient. Nevertheless the system is robust to high personnel turnover, high production demand, and the ever present possibility of major accidents [11]. Ironically the system that encompasses Navy carrier flight-deck operations, where the need for high-levels of production and the ever present possibility of disaster go hand in hand, is reliable and robust specifically because of a mode of organization that members of the Kemeny commission might view as distinctly nonmilitary. This self-organizing mode of operation, in which autonomy for local decisions is transferred to lower levels in the hierarchy, has been a central feature of many successful military operations [7].

INFORMATION-BASED CONTROL

The lesson to be drawn from this example is that effective management will seek to transfer decision making to lower levels of the organizational hierarchy where information is likely to be most detailed and most accurate[2] [12]. The establishment of detailed rules and procedures that constrain operators at the human-machine interface will be counter productive. An effective approach to risk management should encapsulate a form of control in which managers focus on the global state of the system and rely on the experiential knowledge of operators at the interface to detect and to adapt to local fluctuations. Thus, managers and controllors must be sensitive to information that is specific to their own level of influence: global information for managers and local information for operators at the human-machine interface. The primary role of management in this scheme is to establish global constraints such as emphases on safety, timeliness, and productivity.

One implication of this view is that operators at a control interface must be given access to local information about the state and activity of the system and they must be given opportunities to fully assimilate the meaning and importance of that information [12]. Access to information is a problem for design of the human-machine interface; in particular the design of displays. The presentation of inadequate, confusing, and indirect information contributed to the nuclear accident at Three Mile Island [1] and to the destruction of Iran Air Flight 655 by the USS Vincennes [13]. The goal of display design should be to provide operators with direct access to information about the process to be controlled [14]. In particular, operators should not be required to infer or to interpret the meaning of displayed information.[3]

The development of sensitivity to information is a matter of training, although it is training of a special sort. The workup period on a carrier flight deck can be viewed as an opportunity to build and to explore the (information) environment of the workspace. The result is a form of experiential knowledge [7] that permits the flight deck crew to adapt to unanticipated circumstances. More generally, training must provide opportunities for trainee operators to increase their involvement in a full range of tasks as they move gradually toward full participation in the activity. During this process, trainees must become sensitive to local information that specifies system state and critical system behavior. This may be accomplished by on-the-job familiarization of the type described by Rochlin et al. [10], by legitimate peripheral participation in communities of practice as described by Lave and Wenger [15], or by special training scenarios that include simulation of relevant information flows [16].

CONCLUSION

Large-scale industrial systems have the capacity to generate subtle and potentially disastrous interactions. The standard engineering solution to this problem is to create a closed (completely-defined) rule-governed system. In practice, this approach will fail because faults are often unique and because rule sets cannot completely constrain human control action. A rule-based approach to risk management results in brittle systems that are prone to catastrophic failure in the face of unanticipated events. A more robust system can be established by providing controllors direct access to information about critical system states and by permitting them to acquire experience in interpreting the meaning of that information.

[2]The power of adaptive problem solving at the human-machine interface may be illustrated by the anecdotal example of a motorist who, having been wounded by gunfire, used his cellular telephone to seek emergency assistance. He did not, however, know his own location. The emergency dispatcher, on hearing police sirens in the background, was able to locate the motorist by requesting police in the area to switch their sirens on and off. Although the circumstances were unusual, this is hardly a remarkable example of human problem solving. These are the sorts of innovative solutions that human experts can generate. *Of course, this problem solution could be encapsulated* in a rule set. The development of such a rule set does, however, require anticipation of the problem and of the resources available for its solution. This is the ubiquitous problem of control by rules. Is it indeed possible to identify all critical scenarios and relevant possibilities in advance?

[3]The contrast between indirect and direct information was well illustrated in a recent debate between representatives of the food packaging industry and representatives of a consumer group concerned with nutritional value of packaged foods. The issue of concern was how to ascertain whether a product can be characterized as low in fat. From the nutritional perspective, a low fat product is one that has less than 30% of its calories in the form of fat. The view of the packaging industry is that information to calculate this value is listed on food packages and that consumers need to be educated in the use of the required algorithm. The opposing view is that consumers should not be asked to interpret indirect information and that packages should specify the desired value in a direct form that does not require calculation or inference.

REFERENCES

[1] C. Perrow, *Normal accidents*. New York: Basic Books, 1984.

[2] J. Reason, *Human error*. Cambridge, MA: Cambridge University Press, 1990.

[3] H.L. Dreyfus, and S.E. Dreyfus, *Mind over machine*. New York: Free Press, 1986.

[4] G.I. Rochlin, "Iran air flight 655 and the USS Vincennes: Complex, large-scale military systems and the failure of control," *Conference on Large-Scale Technological Systems*. Berkeley, CA, 1990.

[5] L.A. Suchman, *Plans and situated actions: The problem of human machine communication*. Cambridge, MA: Cambridge University Press, 1987.

[6] C. Perrow, "The organizational context of human factors," *Administrative Science Quarterly*, vol. 28, pp. 521-541, 1983.

[7] G.I. Rochlin, "Informal organizational networking as a crisis-avoidance strategy: US naval flight operations as a case study," *Industrial Crisis Quarterly*, vol. 3, pp. 159-176, 1989.

[8] G.B. Rodgers, "Factors contributing to child drownings and near drownings in residential swimming pools," *Human Factors*, vol. 31, pp. 123-132, 1989.

[9] T. Winograd, and F. Flores, *Understanding computers and cognition: A new foundation for design*. Reading, MA: Addison-Wesley Publishing Company, Inc., 1986.

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

[11] G.I. Rochlin, "The case for experiential knowledge," *Second International Workshop on Safety Control and Risk Management*. Karlsbad, Sweden, 1989.

[12] K.H. Roberts, "New challenges in organizational research: High reliability organizations," *Industrial Crisis Quarterly*, vol. 3, pp. 111-125, 1989.

[13] G.A. Klein, "Recognition-primed decisions," in *Advances in man-machine system research*, Vol. 5, W.R. Rouse, Ed. Greenwich, CT: JAI Press, 1989, pp. 47-92.

[14] K.J. Vincente, and J. Rasmussen, "The ecology of human-machine systems II: Mediating 'direct perception' in complex work domains," *Ecological Psychology*, vol. 2, pp. 207-249, 1990.

[15] J. Lave, and E. Wenger, *Situated learning: Legitimate peripheral participation*. Cambridge, MA: Cambridge University Press, 1991.

[16] G. Lintern, "An informational perspective on skill transfer in human-machine systems," *Human Factors*, vol. 33, pp. 251-266, 1991.