NOT FOR QUOTATION
WITHOUT PERMISSION
OF THE AUTHOR

COMMUNICATION AND DECISION AIDS FOR NUCLEAR ACCIDENT MANAGEMENT: PLANNING TO DEAL WITH UNCERTAINTY

John W. Lathrop

June 1981
WP-81-79

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

CONTENTS

COMMUNICATION AND DECISION AIDS FOR NUCLEAR ACCIDENT MANAGEMENT: PLANNING TO DEAL WITH UNCERTAINTY

John W. Lathrop

INTRODUCTION

This article examines some of the communication and decision problems involved in the management of an accident at a nuclear reactor. The words accident management here mean the supervision of the actions taken to mitigate the effects of an accident after it as begun and has been recognized. The purpose of this article is to propose a new approach to aiding communication and decision making in the midst of the confusion and uncertainty that may pervade the next nuclear accident. That approach is concerned with the type of information communicated, not with communication hardware. Specifically, the proposed approach provides a systematic way to encode the information that must be passed from the technical personnel assessing the status of the reactor to the political authorities making decisions concerning population protection countermeasures.

The line of reasoning developed here is based on a particular definition of nuclear accident:

> an occurrence at a nuclear reactor where for some period of time there is a significant probability of an immediate or future release of nuclear material. Significant probability here means a probability high enough that there is or should be consideration of off-site population protection measures, such as evacuation.

This definition is at variance with the common idea of a nuclear accident, which involves some rapid progression of unfortunate events that is understood well enough to allow non-probabilistic projections of off-site doses (see for example IAEA 1979). The probabilistic nature of the above definition reflects the basic idea that motivates this paper: In the course of a nuclear accident, accident management decisions may have to be made on the basis of very incomplete information, so plans should be designed to aid accident managers in making decisions under uncertainty.

The probabilistic nature of the definition used here follows from the experience at Three Mile Island (TMI). Of course, lessons learned from TMI should be treated with great caution, as the next accident will not be the same. Yet the TMI event revealed a type of nuclear accident not anticipated in the emergency plans: a confusing, slowly developing accident with a great deal of uncertainty about current and future plant status in the course of the accident. One of the most significant aspects of the accident at TMI was the very slowly evolving state of information regarding the actual state of the plant (Dieckamp 1981a). This led to long periods of time during which decision makers did not know whether or not a release was about to occur, but only had some unexpressed probability distribution over possible events. Yet accident management decisions could not wait for all uncertainties to be resolved. Decisions had to be made in the face of great uncertainty.

It seems clear from the reports of the Kemeny and Rogovin commissions that existing accident management plans were not adequate to handle the confusion and uncertainty encountered at TMI (Kemeny 1979; Rogovin 1980). Yet why was an accident like TMI not anticipated in those plans? The Nuclear Regulatory Commission (NRC) investigation by the Office of Inspection and Enforcement (NRC 1979a) cites as one reason the fact that emergency planning was geared to those particular types of major accidents where events occur very quickly. Herman Dieckamp, President of the parent company of the utility operating TMI, has pointed out the problems of accident management plans being too narrowly constrained by preconceived notions of what an accident is going to look like (Dieckamp 1981b). That seems to have been the case at TMI.

These points relate to a fundamental problem of planning for nuclear accidents: they are very rare. That fact leads to several difficulties: there is no frequentistic basis for planning, it is difficult to envision the range of different accidents that are possible, and it is even more difficult to assess the adequacy of the accident management plans and preparedness. A more subtle difficulty in planning for such rare events is that past accidents are thoroughly studied, and so are well understood, in retrospect. Plans and preparedness drills may be oriented toward such well-understood accidents, and so are not adequate to deal with an accident that is very poorly understood during its first stages. Once again, that seems to have been the case at TMI.

The central theme of this paper follows from the above paragraphs: Nuclear accident management plans must be made more resilient to the confusion and uncertainty that may be encountered in the course of an actual, poorly understood accident. The remainder of the paper describes one strategy for increasing that resilience.

THE PROBLEM

One of the most important effects of a poor understanding of an accident situation is a lack of summary descriptors of the plant status. Without such descriptors, communication of plant status becomes time-consuming and complicated, and the information-processing load on those involved in accident management is much higher than it would be for a well-understood accident. The Rogovin report (1980) was quite critical of failures to comprehend and communicate at TMI:

> ...the inability of the utility's management to comprehend the severity of the accident and communicate it to the NRC and the public was a serious failure of the company's management.... Moreover, NRC and B&W [Babcock & Wilcox] employees in the control room also did not recognize or communicate critical information. And their offsite organizations did no better, and perhaps worse, than the utility's offsite engineers at GPU [General Public Utilities] in New Jersey in demanding reporting of important information and in rcognizing the significance of the information that they did receive. The...NRC and B&W did no better than Met Ed [Metropolitan Edison]/GPU in reporting critical information up the management chain and acting upon it....[1]

---

[1] Babcock & Wilcox (B&W) is the steam supply vendor for TMI. Metropolitan Edison (Met Ed) is the utility operating TMI. General Public Utilities (GPU) is the parent company of Met Ed, headquartered in New Jersey.

Clearly, communications were a serious problem at TMI, largely because the plant status was not well understood. There were no summary descriptors that could be used to describe the technical situation.

It follows from these considerations that accident management plans should be designed to help participants cope with and convey very uncertain, poorly understood situations. Plans should anticipate and pre-digest as many decisions as possible, so that real-time decision-making resources can be reserved for the completely unexpected aspects of the accident. There are several communication and decision problems involved in accident management, of course. This paper will focus on one of the more intriguing problems: communication of the accident situation from technical people on site to government authorities off site, in support of decisions concerning which population protection measures to employ where. Clearly, this communication and decision problem is vital to successful accident management. Just as clearly, this aspect of the TMI accident left much to be desired, as indicated in the quote from the Rogovin report presented above.

While the basic elements of nuclear accident management are actually rather simple, the decision-making dilemmas that they generate can be very difficult. The best way to structure the on-site/off-site communication and decision problem is to characterize the decisions required. The off-site government authorities must decide which population protection measures, if any, to implement, and where to implement them (out to what radius, in which sector). Possible protective measures, referred to as countermeasures, include advising people to stay indoors (shelter), advising people to leave (evacuation), or issuing radioprotective medication to prevent intake of radio-iodine by the thyroid (IAEA 1979).

The accident situation represents a possibility of radiation exposure to the population. However, because of uncertain plant status, meteorological conditions, etc., typically no one in the midst of an accident can predict with any certainty what population dose will result. On the other hand, the countermeasures themselves entail risk. People can get killed during an evacuation. There are rare negative side effects to radioprotective medication. There are also political and financial costs associated with executing a countermeasure. There are even negative effects to the population caused by the information concerning an accident and the consideration of countermeasures. The Kemeny Commission concluded that the most serious health effect of TMI to the population was the mental distress (Kemeny 1979). It follows that accident management decision making must balance the risks of an uncertain radiation exposure if no countermeasure is ordered with the uncertain costs and risks of ordering a countermeasure which will only reduce, not eliminate, the risk of radiation exposure. The Kemeny Commission conclusion suggests that accident managers should even consider the effects on the public of indecision.

The balancing of costs and risks between ordering and not ordering a countermeasure is only part of the problem of nuclear accident management. A second important part is that no one person can have a full appreciation of all of those costs and risks. It takes people with technical backgrounds and familiarity with the reactor to assess plant status and gain an appreciation of the likelihood of a radiological release. It takes people with legitimate government authority to make the difficult tradeoffs between the costs and benefits of various countermeasures, to order their execution, and accept responsibility for the consequences. Yet the technical people typically do not have government authority, and the governmental authorities cannot be expected to have the technical expertise and familiarity with the plant. In short, knowledge necessary to make responsible accident management decisions is divided between two groups of people. Clearly, good communications between these two groups is paramount. That communication would be difficult enough when an accident is well understood. When an accident is poorly understood, as was the case for the first days at TMI, it becomes extremely challenging to convey the state of knowledge about the situation to the government authorities in such a way that sound accident management decisions can be made.

An example accident management problem may help bring some of the points raised here into focus, and introduce the idea of probabilistic information in on-site/off-site communication. Consider a nuclear reactor sited near a large city. According to the figures of Dr. Jan Beyea in a report made to the New York City Council (see Sugarman 1979), literally thousands of latent cancer deaths in New York City could result from a severe core-melt accident at the Indian Point reactor. Dr. Beyea's figures make it clear that there could be accident scenarios where evacuation of parts of nearby cities could be called for. The long time required to execute an evacuation would mean that it would not be desirable to wait until a large atmospheric release was certain before ordering an evacuation. Yet there would be a perhaps high cost in lives and property for ordering an unnecessary evacuation. There is a very clear need in this case for communicating probabilistic information to government authorities. For example, in the early stages of an accident, with several anomalies in plant status indicators, the operators could judge that there is a 5% chance of a large release, but that no appreciable release could happen for at least two hours. That information could be passed to the off-site government authorities to be used in their determination of whether or not to order an evacuation. However, the use of such information in accident management decisions requires a careful balancing of the societal costs and benefits of evacuating and not evacuating, a process that would best be performed as much as possible in anticipation of an accident, rather than in the stressful and limited time during an accident. For example, the government authorities could have determined in its planning procedures that, given that day's meteorological state, an evacuation would definitely be called for if the probability of large release exceeds 3% with at last two hours warning. It could also have been determined that if that probability falls between 1 and 3%, discretionary factors such as the weather or various political pressures could be allowed to affect the decision between ordering an evacuation and issuing a warning to prepare to seek shelter. Finally, any probability less than 1% would mean that the risks of an evacuation would definitely be greater than the risks of riding out the accident situation with a population prepared to seek shelter, at least for the time being.

The above example illustrates the need for probabilistic information in accident management communication, and the desirability of anticipating how best to react to that probabilistic information. The example does not address the difficulties of generating that information. While several of those difficulties are discussed later, one particular problem in information generation is best discussed now. Note that in the example the 5% probability was enough to call for an evacuation. That means that while an evacuation would definitely be ordered, there is a 95% chance that it would turn out to be unnecessary. Realistically examining the risks and benefits facing the individual operator, is it reasonable to expect him to sound that 5% alarm? Perhaps an operator would view the situation as a 95% chance of ridicule and lost career opportunities. Similar arguments could apply to each individual decision maker in the chain.

This section has described a challenging problem and opportunity for creative engineering: the development of a communication and decision aid system to support the on-site/off-site accident management problem. Such a system should be resilient to the confusion and uncertainty of a nuclear accident. It should be usable by people under a great deal of stress. It should overcome possible resistance to its use by both on-site technical personnel and off-site government authorities. It would be easy to set up the mechanics of such a system. A simple dedicated telephone line would do. The challenge is to develop a communication and decision system to ensure that that phone line is appropriately used. The following sections review the current status of accident management plans regarding this problem, outline possible forms for such a communication and decision system, and describe the necessary next steps in the development of that system.

CURRENT EMERGENCY PLANS

Accident management plans as currently set up typically do not handle confusion and uncertainty well. The IAEA guide for off-site response plans (IAEA 1979) charges the operator with the responsibility of predicting off-site consequences and informing off-site authorities of that prediction without considering problems of effectively communicating the risks of the situation when the operator is very uncertain about the status of the reactor. This fits in with a typical pattern adopted for accident management plans, involving the setting of bounds on measured, anticipated or projected individual radiation doses that could be caused by the accident. Those bounds are to be used as at least partial guidance in recommending countermeasures. A review of papers presented at an international workshop on nuclear accident management reveals several examples of dose-based countermeasure guidelines from European countries (see Clarke and Webb 1981, von Gadow 1981, and Beskrestnov 1981). One such guideline from the U.S. can be found in the Environmental Protection Agency's Protective Action Guides (EPA 1975), which call for mandatory evacuation if dose is to exceed 5 rem whole body (anticipated maximum individual dose). As Clarke points out in the paper cited above, such a guideline represents a balancing of radiation risks versus the risks of the countermeasure itself. However, it is not clear how such a guideline would help in an accident where there is a great deal of uncertainty regarding the status of the reactor. For example, how would such a guideline have helped a decision maker at TMI late on March 30, 1979, who knew only that there was a hydrogen bubble in the reactor vessel, that its flammability was still being calculated/argued, that it might explode and cause major releases of radiation? Where is the anticipated maximum individual dose in that sentence? There is none, though some probability distribution over such a dose is implied. Of course, as it turned out the hydrogen bubble could not have exploded, but the decision maker did not know that at the time. How was he supposed to use the 5 rem guideline when faced with such uncertainty? It could be argued that "maximum individual dose" presumes some maximum accident occurring. But with that logic the 5 rem guideline could be exceeded whenever the plant is in an off-normal mode. Clearly, accident management guidelines should be linked more directly to the information the decision maker is actually apt to have available in the course of an accident.

The Protective Action Guides (PAGs) do leave room for judgment during an accident. For example, while evacuation is deemed mandatory for anticipated maximum individual whole body dose exceeding 5 rem, if that dose is between 1 and 5 rem, ordering an evacuation is subject to the discretion of the government authorities, and can depend on existing constraints (weather, etc.). In this respect, the PAGs are similar to the probabilistic guidance numbers given in the example above: mandatory evacuation if the probability of large release exceeds 3%, discretionary evacuation if that probability lies between 1 and 3%. This parallel forms a very clear contrast between the two concepts of countermeasure guidance: the PAGs are based on anticipated dose, the probabilistic guidance is based on probability of release. The incomplete state of information the operator is apt to have in the course of a poorly understood accident is much more closely represented by the probability number than the anticipated dose number. While it may be difficult for an operator to think in terms of probabilities, that probabilistic information could be transformed into a more usable form, as discussed below.

A different strategy for communicating a poorly understood reactor status can be found in the NRC's Draft Emergency Action Level Guidelines for Nuclear Power Plants (NUREG-0610, NRC 1979b), which effectively stipulates that the operator should call off-site authorities within fifteen minutes if he finds he is confused about the status of the reactor. While that is a step in the right direction, in that it provides for communication even when the operator cannot predict the course of an accident, it does not provide for communication of the operator's uncertain state of knowledge in any structured way. It also raises the problem of the operator's willingness to actually make such a call.

The Emergency Action Level Guidelines referred to above form part of the Criteria for Preparation and Evaluation of Radiological Emergency Response Plans (NUREG-0654, NRC-FEMA 1980). Both the Criteria and Guidelines are very valuable bases for emergency plans, in that they set up a graded scale of emergencies, each calling for a different level of response from off-site authorities (various levels of readiness, advice to take shelter, etc.). However, they still call for the operator to make release and dose projections without giving him a systematic way to communicate a very uncertain reactor status. In addition, these documents call upon the operator to effectively make recommendations concerning what off-site countermeasures to employ where. This may be one way to communicate the operator's estimate of the seriousness of the accident, but it ignores the on-site/off-site division of expertise and authority referred to previously, and does not bear much resemblance to how countermeasure decisions were actually made at TMI, involving many phone calls and consultations between off-site authorities.

The preceding section established the critical importance of communicating the seriousness of a nuclear accident situation from on-site technical experts to off-site government authorities. Accident management systems that base that communication on dose projections or dose-based countermeasure recommendations are not adequate to manage situations where there is important uncertainty about the status of the reactor. Such systems do not handle the problem of portraying an uncertain situation, such as the hydrogen bubble problem at TMI. While such systems could handle a hypothetical, well-understood accident, they are not resilient to confusion, and so fail to deal with problems that can arise in the course of an actual, poorly understood accident.

The NRC and the nuclear industry are correct in continuing their search for better and better indicators of accident status. However, accident planners must acknowledge the fact that it is impossible to anticipate all accidents. Accident managers must be prepared for the fact that the next major accident may be just as confusing as TMI. It follows that emergency plans must include systems to communicate information about uncertain accident situations from technical experts on site to government authorities off site in such a way that sound accident management decisions can be made. Such a system must make communication of probabilistic information feasible in the midst of the confusion and stress of an accident situation; it must be effective in transmitting information from an operator who may face personal incentives to choose a course of action different from the one in the best interest of society. The next section presents one possible basis for such a communication system.

PROBABILISTIC COMMUNICATION SYSTEMS

An Idealized System

This paper deals with communication systems; not with the electronics, but with the type of information communicated. The easiest way to characterize the sort of information called for is to start by describing the state of knowledge that, ideally, should be transmitted from operator to authority. One way to represent that state of knowledge is in terms of subjective probabilities that each of one or two particular dose levels (to the most exposed individual) will be exceeded at each of several different times in the near future. Table I presents an example idealized transmission.

Table I: Probabilistic Information on Near-Future Radiological Hazard

time after present: t	=	.5hr	1hr	2hr	4hr
p($D_{\max} >$ 10 rem)*	=	.005	.01	.01	.05
p($D_{\max} >$ 1 rem)	=	.01	.05	0.5	.10

* Probability that dose to most exposed (off-site) individual will exceed 10 rem, during time period starting with time t and ending with time 2t after present.

The form of transmission represented in Table I is desirable in that it represents the operator's incomplete state of knowledge about the reactor in terms that are most relevant to countermeasure decisions, without cluttering the channel with technical details of no use to the off-site authorities. The probabilities in Table I are the very numbers off-site authorities would need to plug into any formal or informal societal risk/benefit calculations they would be making in weighing whether or not to order an evacuation. Perhaps different dose levels or times would be preferable as row and column descriptors, or perhaps numbers of latent fatalities and thyroid operations would be more relevant measures of radiation loss than dose to the most exposed individual. In any case, this idealized transmission form would involve a table of radiation loss level vs. time, with as cell entries probabilities (or odds) that each radiation loss level will be exceeded at each time. Ideally, the operator or technical experts could express their state of knowledge concerning plant status in terms of this set of probabilities. The government authorities could then use them as bases for their countermeasure decisions. Naturally, the technical people would be given training on generating such sets of probabilities, and the government people would be trained in how to translate such probabilities into guides for their decisions.

The communication scheme just described, simple enough on paper, probably would not work very well in an accident situation, for two very important reasons:

i) in the midst of an accident, where each member of an operator team is under a heavy information-processing load and a great deal of stress, it doubtful that any of them could meaningfully determine a set of subjective probabilities, and

ii) even with the help of extensive training, government authorities cannot be expected to be comfortable and facile enough with subjective probabilities to combine such information appropriately with social and political value information to come to an appropriate countermeasure decision in the midst of the stress and political pressure of a nuclear accident.

A Practical System: the Standard Language System

The problem outlined above with the idealized communication system can be avoided by substituting keywords (or colors, "condition red," or numbers) from a very coarse, pre-determined standard language in place of the sets of subjective probabilities. A keyword would be substituted for each column of subjective probabilities in Table I, or perhaps one keyword would be used in place of the entire table. The standard language would have a limited vocabulary of from three to ten keywords, each associated with a corresponding standard paragraph describing the hazard presented by the accident situation. In the discussion below, the standard language is referred to with the notation $[K_1, \dots K_n]$, where each element K_i is a keyword paired with a paragraph, with higher-subscript values denoting keywords and paragraphs describing more hazardous situations.

The idea of a standard language is not entirely incompatible with the set of emergency action levels (EALs) described in NUREG-0610 (NRC 1979b). That set could be considered a four-element standard language: [Notification of Unusual Event, Alert, Site Emergency, General Emergency]. However, the EAL descriptions lack any information explicitly describing the probability that any particular radiological loss will be exceeded, they lack any information explicitly relating such probabilities to time of occurrence, they are rather coarse in that only the two highest levels involve appreciable off-site risk, and they lack detail in mapping from keyword to countermeasure decision guidance concerning evacuations.

While in the following discussion the language will always be referred to as being used by the operator or operator-team, in a longer-term accident the responsibility for determining the appropriate keywords could shift to technical support staff.

Requirements, Aspects of the Standard Language System

Ideally the Standard Language System (SLS) would induce the operator to move from very coarse aspects of system status to a particular paragraph and keyword. For this purpose, the mapping from system status to paragraph and keyword must be as unambiguous as possible. The operator could then be held liable if, for instance, he transmits a particular plant status, K_2, when it was unambiguously K_3, since post-accident hearings could reasonably determine that he should have used K_3. The SLS also relieves the operator of some of his responsibility, in that he can refer to the mapping guidelines as a justification for his actions. Using the SLS, he is called upon to exercise his judgment less in the stress of an accident, and can benefit from the more considered judgments of the planners who developed the mapping from system status to paragraphs.

Ideally, the SLS would also induce the government authority to act upon the keyword and paragraph received with a particular countermeasure. Given the keyword received and extenuating circumstances, the authority must feel that only one or two countermeasures could be justified in post-accident hearings. The authority is also relieved of some of its responsibility, since it can cite the keyword given as partial justification of the decision. Once again, a mapping from keywords to decision guidelines can be largely developed in the course of emergency planning, so that decision tasks that would have to be performed in the stress of an accident are replaced by more carefully considered judgments made in anticipation of an accident.

Benefits of the Standard Language System

The above requirements call for the SLS to have as unambiguous as possible a mapping from system status to paragraph and a reasonably clear mapping from paragraph and extenuating circumstances to particular countermeasure. This raises the possibility of a more direct language, mapping straight from system status to countermeasure. Yet the intermediate nature for the SLS presented above is preferable to a more direct language, because it divides responsibility between the operator and the government authority in an advantageous way. This division of labor by the intermediate SLS leaves the operator with a relatively technical judgment and thus less prone to bias his estimate in the light of potential non-technical consequences. Conversely, the government authority, left with a relatively non-technical judgment, is able to bring in other considerations (politics, weather) in a cleanly structured way.

The SLS allows the structured "pre-digestion" of very difficult judgments and decisions that could not be made well in the heat of an accident. The net effect would be to decrease hurried, individual human accident management judgments made under stress, replacing them with more carefully considered judgments made by larger numbers of people over longer periods of time.

The SLS would provide an appropriate avenue for openness with government and the media, and so would help prevent the loss of credibility that so complicated the TMI accident (see Kemeny 1979, Rogovin 1980). In sum, the SLS would build into the accident management system a pre-determined means of describing and reacting to a very uncertain current and future plant status.

Difficulties in Developing the Standard Language System

Some of the major difficulties to be encountered in the development of the SLS are as follows:

1. Operator Stress: The operator will be under severe stress in the course of an accident, so the mapping from system parameters to paragraphs and keywords must be kept simple, and must use only coarse system parameters.
2. Novel Accidents: No new accident is like any past accident. The TMI accident, for example, was quite novel. As a result, the predetermined accident alarm levels were not effective in providing timely indications of hazard. It follows that the mapping from system parameters to paragraphs and keywords must be kept general, and should not be too constrained by preconceived notions of what an accident will look like.
3. Size of Language Vocabulary: The choice of the number of paragraphs and keywords to include in the SLS involves a difficult balance: the greater the number of paragraphs, the more guidance given to the government authority. But as the number of paragraphs increases, the more difficult will be the operator's task of choosing among them.

4. Phrases Used in the Paragraphs: The wording of the paragraphs determines the relative difficulty of mapping from system parameters to paragraph versus mapping from paragraph to countermeasure decision. Paragraphs worded in system status terms are easy for the operator to select, but are difficult for the government authority to use. Alternatively, paragraphs worded in countermeasure terms are difficult for the operator to select, but are easy for the authority to use. Ideally, the paragraphs should be worded in terms of present and future hazard, capturing the information contained in Table 1, and dividing the judgments required as cleanly as possible: technical judgments in mapping from system status to paragraph, political and social judgments in mapping from paragraph to countermeasure decision.

5. Paragraph to Decision Guide Mapping: The mapping from standard paragraph to countermeasure decision guides cannot be determined by the technical experts alone. The appropriate government authorities also have a responsibility to participate in that determination since such a mapping must involve social value judgments that only they can make. These guidelines should be periodically thought through with the relevant government authorities in the course of maintaining preparedness for an accident.

STEPS IN THE DEVELOPMENT OF AN SLS

The development of an SLS represents a very challenging, fascinating set of engineering problems. While there are several ways one could go about such a development, one in particular will be described here as a way of illustrating the engineering problems involved. The first three of the six steps described below involve the identification of classes of radiological loss vs time arrays such as those presented in Table I. These arrays will be referred to as loss-time arrays. While the particular levels of radiological loss and time used are subject to adjustment in the course of the following steps, once the most appropriate levels are determined those levels, as row and column descriptors, will be the same for all arrays. That is, loss-time arrays will only differ in the cell entries, the probabilities that each loss will be exceeded at each time. Six steps in the development of an SLS are now briefly described.

1. Identification of classes of loss-time arrays discriminable by coarse system parameters.

Once the set of all possible loss-time arrays is identified, it is to be partitioned into classes in such a way that an operator looking at very coarse, general system parameters could most easily tell which class includes the array that best describes the hazard presented by the reactor. Another way to describe this task is the partitioning of the loss-time arrays into classes such that coarse, general system parameters are most apt to identify the class of arrays that includes the array that best describes the hazard presented by the reactor.

This step can be very challenging, as it involves the characterization of the reactor system by that set of system parameters that are directly observable and are apt to catch all significant accidents, even very novel ones that have not yet been imagined.

2. Identification of classes of loss-time arrays that discriminate among countermeasure decision guides

A partition of loss-time arrays on the criterion of Step 1 alone would not necessarily lead to a useful set of classes of arrays. The classes of arrays ultimately sought are to form the bases for SLS paragraphs. As such they must discriminate among alternative countermeasure decision guides. In this Step 2, then, the set of all possible loss-time arrays is to be partitioned into classes in such a way that a government authority looking at any of the classes could most easily tell which countermeasure decision guide is the most appropriate for the hazard represented by that class of arrays. Another way to describe this task is the partitioning of the loss-time arrays into classes such that each class clearly identifies a most appropriate countermeasure decision guide.

3. Selection of classes of loss-time arrays appropriate as bases for SLS standard paragraphs

The sets of classes identified in the previous two steps can now be compared to select a set of classes of arrays that combine the characteristics sought in each of the two steps. That is, a set of classes of arrays can now be identified where each class is discriminable from the others by coarse reactor system parameters, and each class in turn identifies a different appropriate countermeasure decision guide. In addition, it is hoped that the arrays within each of the identified classes have enough in common with each other that Step 4 is feasible.

4. Develop an english language paragraph and keyword for each of the array classes identified in Step 3

These paragraphs should form a set that retains the characteristics of the set of array classes identified in Step 3. That is, each paragraph should be discriminable from the others on the basis of coarse reactor system parameters, and each paragraph should in turn provide unique guidance as to which countermeasure alternatives are appropriate. The particular language used in these paragraphs should be selected in accordance with the considerations listed in point four of the list of difficulties presented above.

5. Develop clear guidelines for mapping from reactor system status to each of the standard paragraphs developed in Step 4.

6. Develop clear guidelines for mapping from each of the standard paragraphs to countermeasure decision guidance.

CONCLUSION

This article has focused on a particular problem of nuclear accident management: the communication of information concerning the radiological hazard of a poorly understood accident from technical people on site to governmental authorities off site in such a way that sound countermeasure decisions can be made. Current accident management plans do not handle that problem well, in that they implicitly assume that the accident is reasonably well understood, at least understood well enough for the operator to provide non-probabilistic radiological release and dose estimates. Such plans are not resilient to the confusion and uncertainty that may accompany a nuclear accident such as the one at TMI.

A particular form for on-site/off-site communication of a poorly understood accident situation has been proposed here, called a standardized language system (SLS). The mechanics of such a system involve simply a dedicated phone line between technical experts and government authorities. The development of such a system would concentrate on the keywords passed down that phone line, and the mappings from system status to keyword and from keyword to countermeasure decision guidance. The value of such a system lies in its ability to encode the highly uncertain off-site radiological hazard represented by a poorly understood accident in a form that is usable by the parties at either end of the line.

The development of a standardized language system would represent a challenge and opportunity for the engineers involved. It calls for creative thinking concerning the behavior of a complex system operating in an extremely rare mode. The SLS involves very directly and immediately the interaction between a complex technology and society, between technologists and government authorities. Finally, the development of an SLS could contribute substantially to the safety of nuclear power.

REFERENCES

Beskrestnov, N.V. 1981. Organizational aspects of control and management during emergencies at nuclear power plants in the USSR, in J.W. Lathrop (ed.), *Planning for Rare Events: Nuclear Accident Preparedness and Management.* Oxford: Pergamon Press.

Clarke, R.H., and G.A.M. Webb. 1981. Radiological protection criteria for controlling doses to the public in the event of unplanned releases of radioactivity, in J.W. Lathrop (ed.), *Planning for Rare Events: Nuclear Accident Preparedness and Management.* Oxford: Pergamon Press.

Dieckamp, H. 1981a. The accident at Three Mile Island: From the perspective of the owner operator, in J.W. Lathrop (ed.), *Planning for Rare Events: Nuclear Accident Preparedness and Management.* Oxford: Pergamon Press.

Dieckamp, H. 1981b. Quote from discussions, in J.W. Lathrop (ed.), *Planning for Rare Events: Nuclear Accident Preparedness and Management.* Oxford: Pergamon Press.

EPA. 1975. *Manual of Protective Action Guides and Protective Actions for Nuclear Incidents.* Washington, D.C.: Environmental Protection Agency.

IAEA. 1979. Planning for Off-Site Response to Radiation Accidents in Nuclear Facilities. IAEA-TECDOC-225, Vienna, Austria: International Atomic Energy Agency.

Kemeny. 1979. *The Need for Change: The Legacy of TMI.* Report of the President's Commission on the Accident at Three Mile Island, John G. Kemeny, Chairman. Oxford: Pergamon Press.

NRC. 1979a. *Investigation into the March 28, 1979, Three Mile Island Accident by Office of Inspection and Enforcement.* NUREG-0600. Washington, D.C.: Nuclear Regulatory Commission.

NRC. 1979b. *Draft Emergency Acton Level Guidelines for Nuclear Power Plants.* NUREG-0610. Washington, D.C.: Nuclear Regulatory Commission.

NRC-FEMA. 1980. *Criteria for Preparation and Evaluation of Radiological Emergency Response Plans and Preparedness in Support of Nuclear Power Plants.* Nuclear Regulatory Commission and Federal Emergency Management Agency. NUREG-0654. FEMA-REP-1, Rev. 1. Washington, D.C.: Nuclear Regulatory Commission.

Rogovin. 1980. *Three Mile Island.* Report of the Nuclear Regulatory Commission Special Inquiry Group. Mitchell Rogovin, Director. Washington, D.C.: U.S. Nuclear Regulatory Commission.

Sugarman, R. 1979. Nuclear power and the public risk. *IEEE Spectrum 16* (11), November.

von Gadow, A. 1981. Emergency planning in the vicinity of nuclear installations, in J.W. Lathrop (ed.), *Planning for Rare Events: Nuclear Accident Preparedness and Management.* Oxford: Pergamon Press.