Reliability Engineering and System Safety 36 (1992) 207–216

Cognitive task analysis of a complex work domain: a case study

Neville Moray, Penelope M. Sanderson & Kim J. Vicente

Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

Conducting an analysis of a work domain is a necessary prerequisite for making sound and effective human factors design recommendations. There are several properties of complex work domains which put very strong requirements on the type of analysis methods which can be meaningfully used. Traditional task analysis methods based on observable actions are no longer appropriate. In this paper, we describe a set of methods that we have adopted and have found to be useful in conducting a cognitive task analysis of a complex work domain. From our experience with this particular case study, we have found that it is necessary to adopt a variety of different methods to examine a particular work domain, and use the converging evidence from these several methods to attain a final, comprehensive understanding of that domain.

INTRODUCTION

In recent years there have been two important developments relevant to the analysis of human–machine systems. The first resulted from the increasing introduction of automation into the workplace, which caused a change in the nature of the tasks which faced operators. Increasingly, it has become impossible to define work in terms of physical events, since much of the work has changed from physical manipulation to mental decision making. The result has been the introduction, or the attempt to introduce, methods of *cognitive task analysis* to replace or supplement traditional task analysis (e.g. Rasmussen;[1] Woods and Hollnagel[2]). The second development has been an increasing emphasis on the analysis of groups or teams. Very seldom is it the case that a task is performed by an isolated operator. Some very successful analyses of team performance have appeared (e.g. De Keyser *et al.*;[3] Rochlin *et al.*[4]). But if cognitive task analysis has been found to be difficult because of the relatively small proportion of overt observable behavior, the problem of combining cognitive task analysis with the analysis of group and team behavior is even more beset with difficulties.

In this paper we wish to report our experience of trying to perform such an analysis. We will not be able to report a successful methodology in its complete form, since the work is still in progress. But we will indicate some of the approaches which we have found to be helpful. If there is one theme which we would emphasize, it is that it is necessary to adopt a variety of methods to examine a particular work domain, and use the converging evidence from these several methods to attain a final, comprehensive understanding of that domain.

We also wish to explore the relation between detailed analysis of tasks as they are performed and the use of general models of the work domain. A direct attack on the details of performance, as is typically adopted in classical task analysis methods (Meister[5]), may leave one with too contingent an analysis, where there is an excessive dependence on the particular details of activity which one has seen. The development of a general model overcomes this limitation by providing a powerful aid to thinking about the entire range of demands confronting the operators, thereby allowing one to evaluate system functionality (i.e. does the structure of the domain and the available information make it possible to achieve the desired objectives?). On the other hand, such a global framework does not address some important detailed design issues (e.g. room layout). More traditional and more specific human factors methods are required for these types of problems. Clearly, both classes of methods are required for a comprehensive analysis.

Reliability Engineering and System Safety 0951-8320/92/$05.00
© 1992 Elsevier Science Publishers Ltd, England.

A CASE STUDY: AN EMERGENCY MANAGEMENT CENTER

The facility in which we have been working is an emergency management Center. The task of the Center is to monitor the status of several nuclear power stations within a radius of some 150 miles, and, in the event of an accident, to make decisions as to the nature of the problem and the probability that there will be a release of radioactivity which will be dangerous to the public. If the latter is deemed probable, or if it in fact occurs, the responsibility of the Center is to recommend an appropriate response strategy. Should a section of the population be asked to shelter, should it be evacuated, or is neither necessary? How do meteorological conditions determine which areas are at risk, and how does the status of the reactor interact with expected environmental conditions to determine what the course of the emergency will be over a period of hours or even days? The actual execution of the recommended response strategy is carried out by another agency.

The people who man the Center are highly qualified, well trained, and highly motivated. Most have had prior experience in connection with nuclear power, either in its civilian or military manifestations. Regular exercises are held, and plans are continuously made and implemented to upgrade the efficiency of the Center. It is in connection with these efforts that we have been working on the cognitive task analysis.

Data are constantly received at the Center from all of the reactors over dedicated lines which are sensed and processed by computers, so as to supply real time monitoring from the same sensors which are monitored at the reactor sites. In addition, similar communications networks provide access to meteorological and radiological data, which can be supplemented by measurements made by field teams and sent to the Center either by phone, fax, or computer link.

There are about 15 people in the Center during a full-scale incident, including a team which analyses the data received from the reactors, a team which analyses the environmental data, both from a meteorological and health physics point of view, a computer specialist, and two executive officers. The reactor team and the environmental analysis team are each composed of a team leader and three assistants. There are, in addition, support staff who handle incoming radio messages, incoming phone calls from the Press, the State authorities, Regulatory agencies, etc. In the event of a severe incident, the number of people in the Center can be increased by about five or six as people arrive whose specific responsibility is to provide liaison with outside agencies. The Center is supplied with computer, fax, phone and radio links with outside agencies, and it is to such agencies that the Commander of the Center directs his decisions as to what actions should be taken.

We have observed two full scale exercises to date. These last for many hours and involve an extremely realistic pattern of events which are provided by a powerful simulator system. During such exercises the utility which owns the power reactor is also exercised using its training simulator, and exactly the same data appear at both sites. At the same time the field teams, police, and other emergency response agencies are also involved. We have been very impressed by the efficiency and dedication of all those involved.

THE PROBLEM: TEAM-ORIENTED COGNITIVE TASK ANALYSIS

We have tried to understand the relation between the actions of this highly coordinated team operating in real time to analyze highly complex, uncertain, dynamic data. Our hope is to detect patterns in the behavior of the system which indicate points of weakness or potential overload, and to then be able to make recommendations for redesign to support more effective operation.

It is clear that a conventional task analysis will not suffice for two reasons. First, the overt behavior consists largely of watching computer displays, typing in commands, and communicating over the phone. To look merely at overt behavior is to miss most of the significant events, since the major work of the team is to assess information and to make decisions. The results of these activities appear as overt behavior, but the nature of the process cannot be traced in that way. Furthermore, the essence of the operation is an intense interaction both within and between the subgroups (reactor analysts, RA; environmental analysts, EA; and the executive officers, XO). Because of the inherent complexity of the operations of the Center, it is only possible to sample a very small range of its activities. The exact behavior exhibited will vary widely from one exercise to another, depending on: the scenario which is chosen (the team is not told anything about the exercise in advance other than which utility is involved); the current configuration of the Center; and which personnel are currently manning the Center. Furthermore, it is extremely difficult to collect real time data on the interactions of 15 people working on real time dynamic problems in a room approximately $10\ m^2$.

How can one arrive at a valid description of the activities of such a team involved in intense, dynamic cognitive work?

A TENTATIVE SOLUTION: CONVERGING EVIDENCE

We believe that it is necessary to draw on a variety of methods and sources of data which will, when combined, provide an integrated description of the operation of the Center which has plausible face validity. It will then be possible to validate that description by predicting critical events, and by comparing the pattern of operations observed during exercises with that expected from the results of the analysis. In this paper, we will describe the first part of this process, reserving an account of the predictive validation for another occasion.

We have made use of four approaches, namely:

1. observational field studies of exercises;
2. analysis of normative procedures;
3. questionnaire and interview data;
4. a formal model to provide a framework for combining the first three, and to provide a means of generalizing the results.

1 Observational field studies

During exercises, video recordings and audio recordings were used to provide a general impression of the events over a period of hours. In addition, three observers were used to monitor and time sample the events. Each observer was responsible for recording data from one of the three groups: RA, EA and XO. Using synchronized timing, they noted at 5-min intervals what each member of their subteam was doing, what equipment was being used, with whom communication was taking place, and, as far as possible, what the content of the communication was at the time. These data were transcribed into timeline protocols, and were compared with one another to try to establish gross patterns of interaction, what the main choke points were, etc. The protocols were checked against the audio and video recordings. (These latter were not very helpful except when used to clarify the transcribed timeline protocols, because of the problem of visual and audio crosstalk in a very crowded room where only two cameras could be used.) An example of a section of the protocol is shown in Fig. 1.

A preliminary account of the dynamics of the interaction between personnel was obtained from the timeline protocols using Hypercard. At each 5-min point, a complete list was obtained from the protocol of who was talking to whom, how they were moving about the room, and what equipment they were using. A Hypercard stack was prepared which showed the layout of the Center complete with equipment (phones, computers, etc.) Each page of the stack was filled in with a graphical representation of the interaction at the particular 5-min sample point, using arrows to indicate movement, interaction and communication. When this stack is played back, a 'movie' is obtained showing the way in which events unfold in time. The animated sequence is quite effective in revealing moments when there is overload, or patterns of inefficient movement and congestion. In effect, it provides an animated link diagram. Two frames are shown in Fig. 2.

This representation can provide useful information for discussing possible changes in the layout of the Center to minimize the time spent in 'travel', and for identifying places where equipment could be rearranged or where extra facilities are needed.

At the end of this phase of the analysis, we had obtained evidence of the patterns of interaction between members of the teams, patterns of equipment utilization, and the ways in which information was obtained and decisions communicated, both within the Center and to agencies outside the Center.

At a later stage in the project, the use of Hypercard stacks turned out to be of value when presenting design changes. On the basis of the analysis we concluded that a change in the location of the Commander's office would be desirable. Instead of his room being off to 'top right hand corner' (as seen on the Hypercard diagram) of the main room, we recommended that the wall on the 'bottom' of Fig. 2 should be removed, and the Commander's office placed so that it looked into the room from that location. Doors should be placed at both ends, opening into the main room, so that the reactor analyst and the environmental analyst could each reach him directly without having to cross the main room and negotiate tables and other congested areas. The new layout was printed on Hypercards, and the activities which we had observed in the initial analysis modified to reflect what would have happened had the layout been as we suggested. This provided a rapidly created and dynamic representation of the savings and simplification of movement which would be obtained in the new layout, see Fig. 3.

12:55			
	null		
13:00			
	REA	-	asked by REAC Executive Officer what maximum projected dose is.
		-	gives him an answer
		-	goes into meeting with REAC Commander
	REAA	-	giving status report on phone
	RECS	-	tells EC of new WD and WS
	EC	-	receives information from RECS and updates board
13:05			
	null		
13:10			
	REA	-	receives phone call via REAC Communications dispatcher regarding farms
		-	tells REAA to calculate dose rates
		-	asks REAA what noble gas rate is
	REAA	-	receives command and question from REA
	RECS	-	null
	EC	-	null
13:15			
	REA	-	on phone to Iowa Liason
	REAA/RECS	-	working on PC projections
	EC	-	updating status board

Fig. 1. Section of time line protocol. (REA, Environmental Analyst; REAA, Environmental Analyst Assistant; EC, Environmental Communications Operator; RECS, Reactor Communications Operator.)

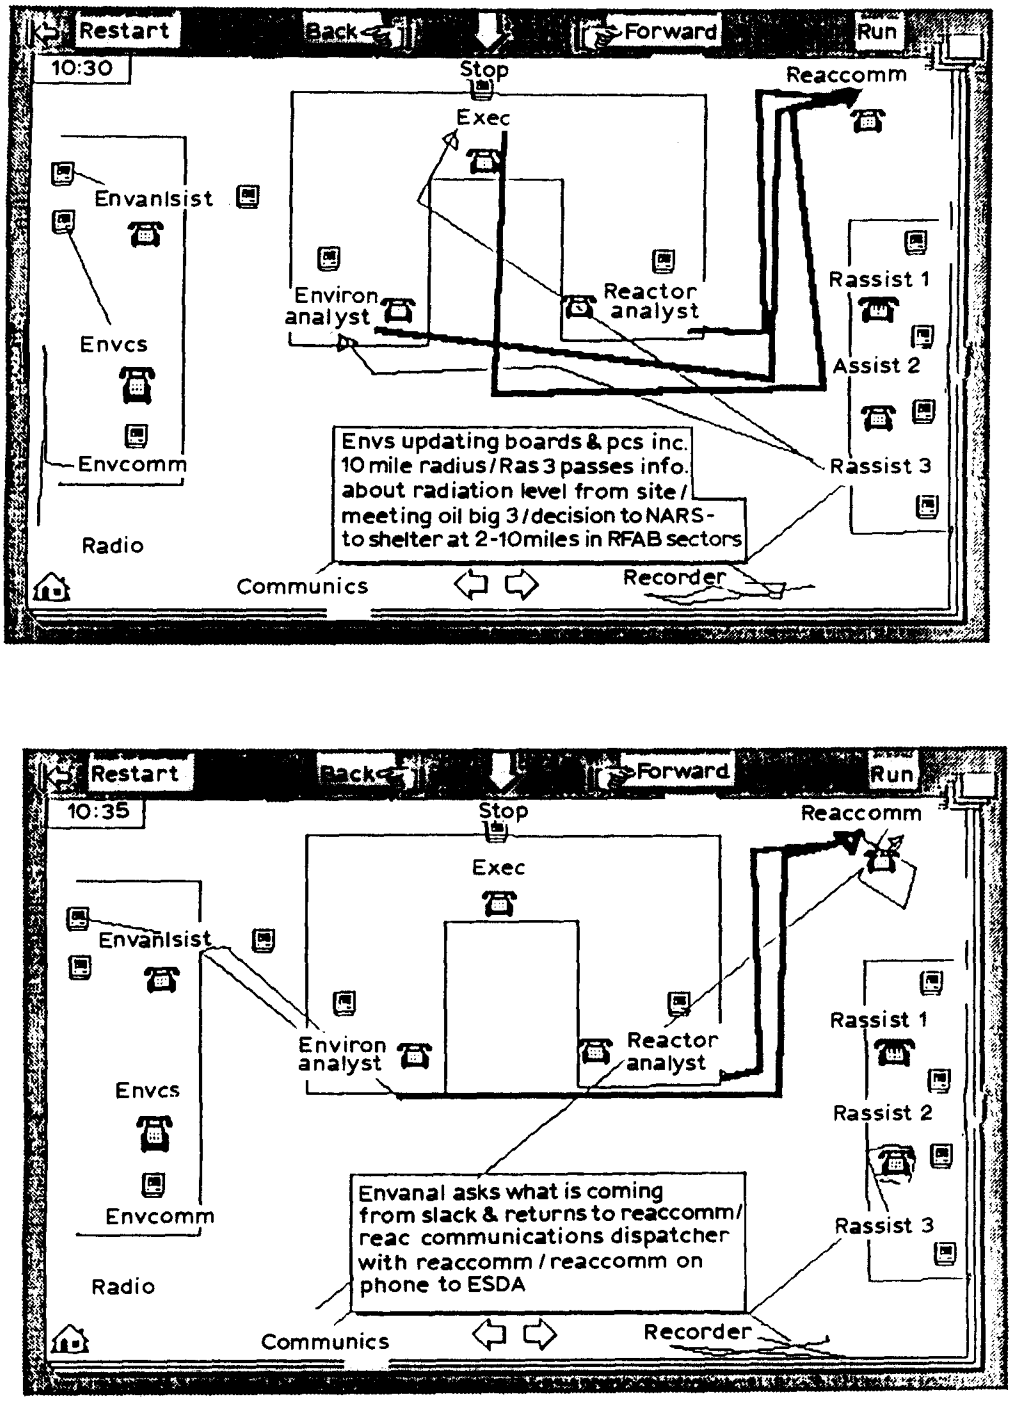

Fig. 2. Sequential frames from Hypercard stack showing movements and activities in control room original configuration. Heavy lines show movements of people. Light lines show electronic communication or voice communication.

2 Analysis of normative procedures

In any facility which has a well defined mission, formal procedures exist. These take the form of procedures which must be followed, documents which define the allocation of responsibility within the facility and its relations with other agencies, and so on. A close examination of such documents will reveal the pattern of behavior which the facility is meant to show in all foreseen cases. Most of these prescriptions will consist of what we may call 'behavioral responsibilities'; that is, things such as 'Notify Agricultural Authorities in the region if livestock must be put on stored feed'; 'Determine probable time and pathway of release of radioactivity'; and so on. These are quite specific actions, and it is easy to notice their occurrence.

But underlying these prescribed behaviors there is a more subtle set of implications. If the final outcome of the Center's activities is to make such decisions and to

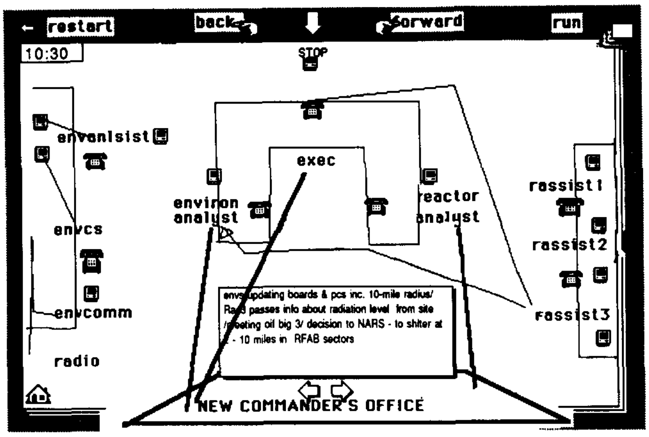

Fig. 3. Frame from Hypercard stack showing pattern of movement predicted with new layout. Note that the commander's room has been moved from the top right corner to the bottom of the figure. Compare with card for '10:30' in Fig. 2.

issue appropriate commands, then certain cognitive activities will be required. In order, for example, to determine whether a release is probable, and to predict its time, the RA team must observe the values of certain sensors, monitor the time course of certain reactor parameters, diagnose the probable nature of the incident and predict its likely course, etc. It is not possible to arrive at the final decision without these steps. Similarly, on the EA side, to decide in what area of the country people or animals should be sheltered or evacuated requires observations of the meteorological conditions, predictions of future meteorological conditions on the basis of these data, deductions about the likely concentration at ground level of radioisotopes of certain kinds in certain areas, etc. Again, in order to arrive at the prescribed conclusions certain cognitive activities must, logically, be performed. Hence, in each case, it is possible to deduce the cognitive components of the task which are logically implied by the formal normative procedures defined by the agency for which the facility works.

Such an analysis of the documents which define the mission of the facility can therefore be used as a second source of information for constructing a cognitive task analysis. Typical of the kind of document which such an analysis produces is a decision flow chart which summarizes at a high level the major choice points which are involved in the decisions which have to be made. Such a flowchart is shown in Fig. 4.

The analysis of the normative prescriptions immediately raises a number of questions, relevant to the identification of any deficiencies in the facility. Is the list of prescribed tasks adequate in principle to cover all important incidents? If so, does the design and operation of the facility support all the requirements of the mission prescriptions, at least in principle? If so, do people actually carry out the tasks which are mandated in the way they are mandated? If not, is the failure to do so the result of poor facility design and layout, poor equipment, inadequate training, or some more subtle factor such as the occurrence of momentary overload of operators due to the concatenation of cognitive demands under certain circumstances?[4] Do coordination and communication occur in efficient patterns, and in patterns which satisfy the requirements of the decision points indicated in a graph such as that in Fig. 4? That is, more generally, does the actual behavior of the operators match the prescribed actions required to fulfill their mission?

3 Questionnaire and interview data

The field study data, timeline protocols and Hypercard animation are very task specific. If a different incident had been observed, or the same incident at a different utility, it is likely that the exact details of those records would have differed substantially from what was observed. On the other hand, the normative procedures are generic, and are intended to apply to all incidents. We have used a semistructured interview and questionnaire technique to bridge the gap between these two. The aim here is to see the tasks through the eyes of the operators.

For each decision node of Fig. 4, the personnel of the Center were asked a series of questions. When that particular decision was being made (collectively by the entire Center), with whom did they interact? Where did they get information? What equipment was used? To whom did they pass information or commands? What were the main sources of difficulty or workload?

The resulting data were tabulated for all personnel for each decision point. The aim here was to build up a summary representation of the kinds of data which were represented in detail in the animated Hypercard stack. The method of presentation chosen was an Interaction Matrix, as shown in Fig. 5.

In these matrices, the rows and columns of the matrix represent particular generic personnel, such as the Reactor Analyst, Environmental Analyst Assistant, Executive Officer, etc. A cell which is shaded indicates that the person on the row margin reported that he or she interacted with the person on the column margin. The hatched cells indicate implied interactions. For example, in some cases people did not say that they interacted with another person, but did, in another question, state that they received information from the latter. In such cases there is an implied interaction which is represented by a hatched square. (It is interesting that there seems to be a tendency for people to interpret 'interaction' as meaning that they give information or commands to people. If all they do is receive information, some at least do not describe this as an interaction.)

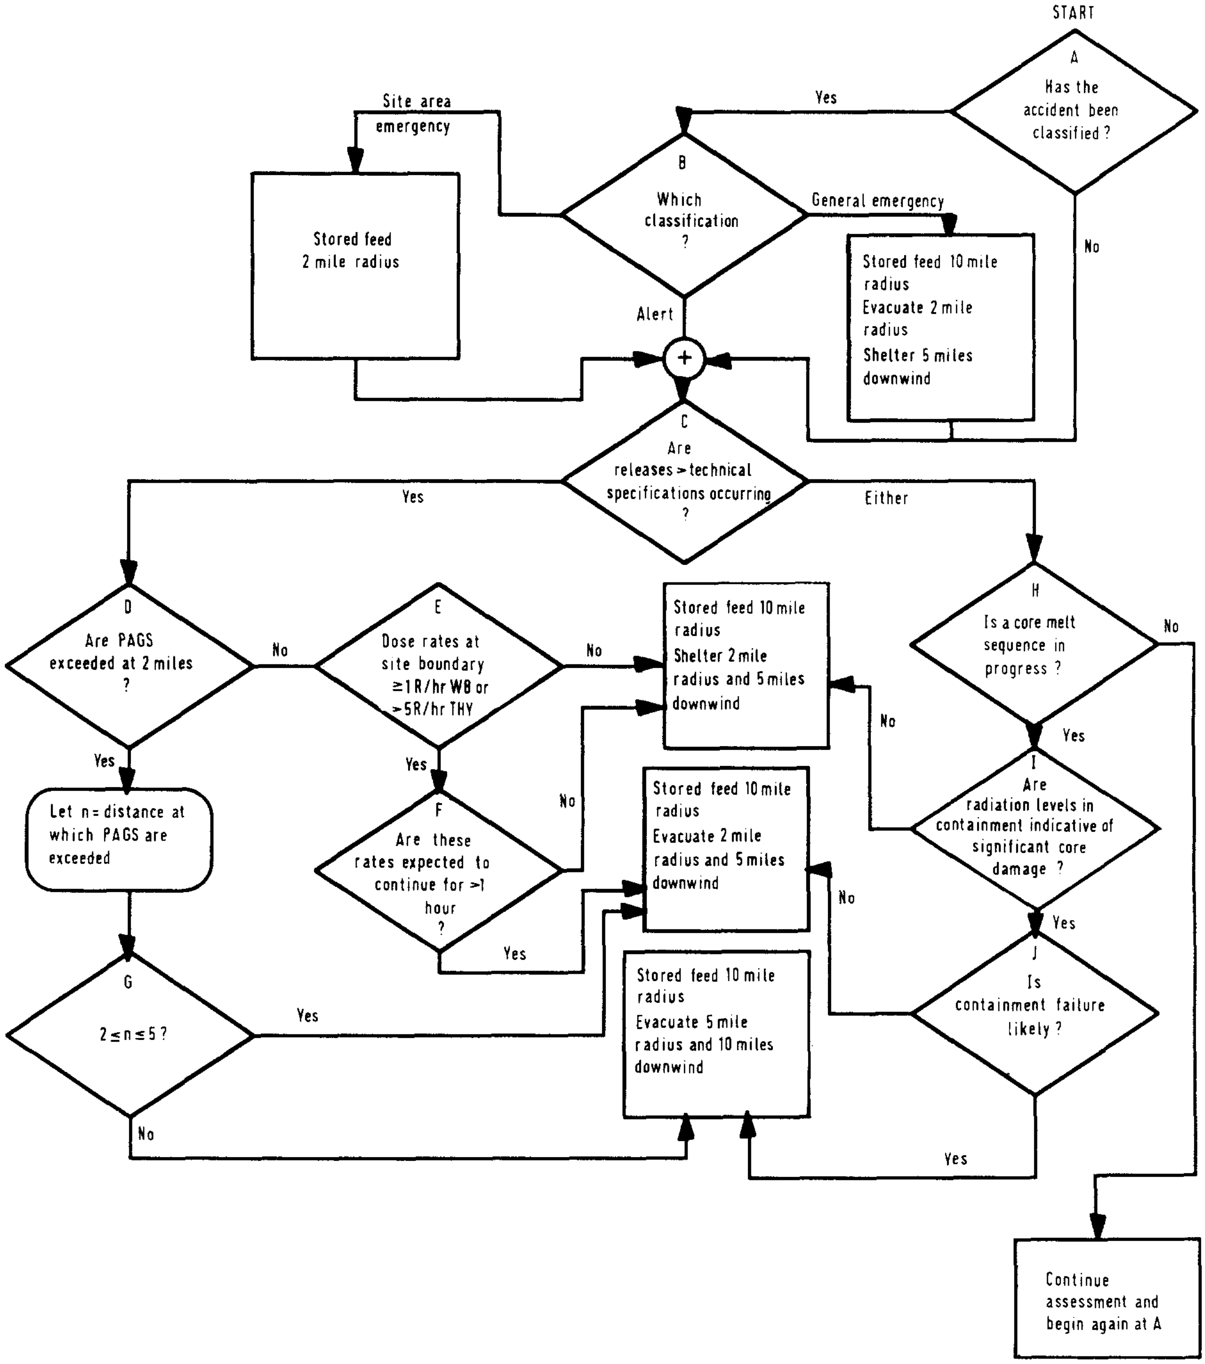

Fig. 4. Decision tree derived from official response plan. One aim of the analysis is to establish to what extent this plan is actually used, and to what extent other patterns have been developed during practical experience. Operating procedures assume the diagram is accurate.

In addition, the outer left and right columns of the matrix are used to show sources of information and destinations of messages. Without going here into the specific interpretation of what these refer to in the case we are investigating, clear patterns emerge which indicate that certain people have a very heavy workload. Again, such information can be used for design recommendations.

These matrices capture quite a lot of the cognitive task loading of the operations of the Center. In several cases, they clearly show differences from the implications of the decision tree as shown in Fig. 4. The most striking feature here is the fact that there seems to be a periodic yet critical interaction between the two team leaders (RA and EA), often mediated by the XO. This disagrees with the almost complete separation of function implied by Fig. 4. To discover the existence of such interaction is of great importance in planning changes in layout and organization of the Center, since it shows that there are important

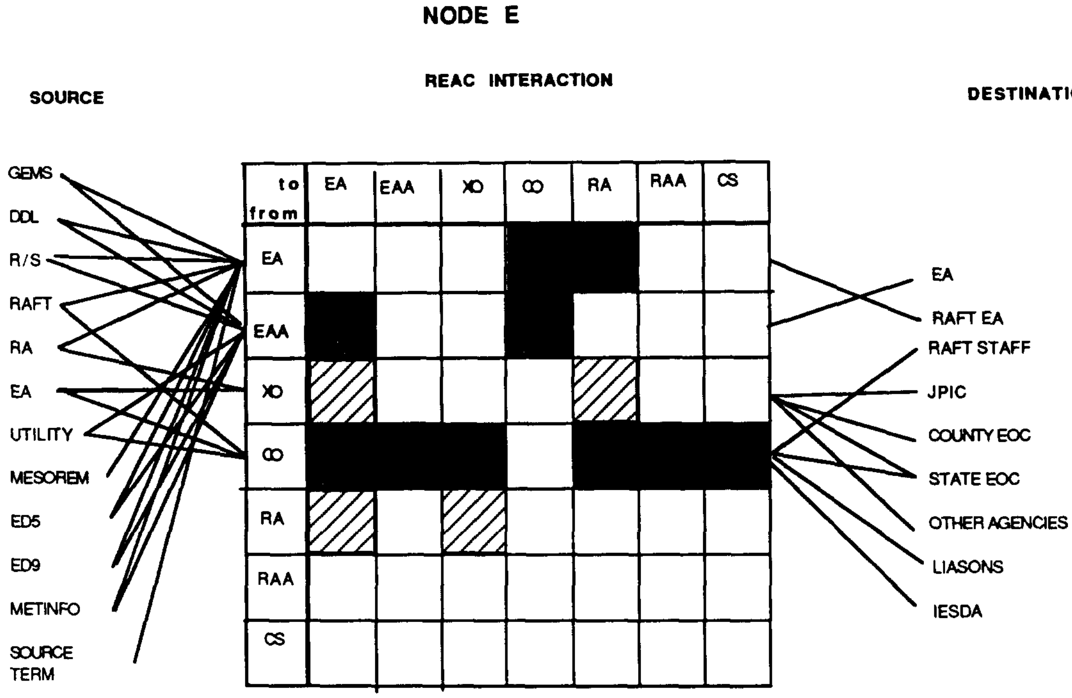

Fig. 5. Matrix of interactions among personnel at decision node E of Fig. 4. Dark squares are interactions reported during interviews. Hatched squares are interactions implied by comments but not expressly identified. This figure could be converted to a three-dimensional grid in which equipment used in each interaction would be represented in the third dimension.

interactions which must be supported at all times, even though there does not seem, in principle, to be the need for this is in the normative plan.

Another value of these matrices is in providing a focus for feedback and discussions with personnel. When shown the matrices most personnel said that they did not represent an accurate picture of the interactions which took place, despite the fact that the matrices were based on interview and questionnaire data. Subsequent discussion allowed the patterns of interaction to be established with much greater accuracy, and the matrices played a central role in those discussions.

4 A formal model

As mentioned earlier, in facilities such as these it is impossible to sample more than a very small proportion of the behaviors which might be seen. The Center monitors more than ten reactors, with several thousand data points coming into the computers every few minutes. The number of potential incidents which could in principle occur is enormous, and the meteorological conditions virtually infinite. Even if the methods outlined above could provide us with a reasonable cognitive task analysis of the behaviors and interactions which we have examined, how can we generalize to this vast population of potential activity? Clearly, the only feasible way is to try to develop a generic model that is based on the broad structure of the domain demands, rather than on the observable behavior unique to particular incidents. We believe that such a model can be derived from the framework suggested by Rasmussen.[1,6,7]

The general model we have adopted consists of three dimensions:

1. A means-end hierarchy representation of the demands of the work domain that is independent of the details of any particular incident.
2. A high level description of the decision activities that personnel are faced with, described in terms of the decision ladder.
3. A representation of the work organization describing how the different decision tasks are allocated among people, and how the activities are coordinated.

The three dimensions of the model correspond to three fundamental questions that a comprehensive cognitive task analysis must provide answers to. These are:

1. What is the functional structure of the work domain?
2. What are the decision activities associated with the problem?
3. Who is responsible for performing the various decisions?

Due to space restrictions, we will concentrate here on describing the means-end hierarchy representation of the work domain.

The global work domain of which the Center is a part is presented in Table 1, and is composed of two parts which are relatively independent. The Domain of Potential Risk describes potential sources of accidents, and the potential effects of accidents. One can think of this part of the domain as a dynamically changing set of accident sources and potential effects arranged in a geographical map. The Domain of Mitigation Resources, on the other hand, describes the range of accident control which may be called into use in case of an emergency. It is important to note that the functions outlined in Table 1 should represent the entire range of possibilities. Typically, for any specific incident, only a subset of those functions will be activated and under consideration. In what follows, we will describe the Domain of Potential Risk which is more central to the activities being conducted by the Center.

The Domain of Potential Risk describes the way in which the effects of accidents can propagate, as well as the potential higher level consequences that can occur. As shown in Table 1, there are five levels of representation. We will describe each level in turn, starting at the bottom of the hierarchy.

The level of physical form is the most detailed of all. Here, the domain is represented in terms of physical attributes (e.g. location and appearance). This would include a description of the different systems, subsystems, and elements of each reactor (as would be found in blueprints, for instance). This level would also include geographical representations of the population density around each site and the prevailing meteorological conditions. Since it is impossible to anticipate the particulars of an incident, this level of description must include all of the reactors and sites in the state. When an accident occurs, only the information associated with the affected reactor would be used.

At the next level, physical function, the physical entities described at the lower level are described in terms of their functional capabilities. This would include the different ways in which the weather can interact with the accident characteristics in order to affect the consequences of a release. In terms of the reactor, we need to know the different accident mechanisms which can lead to a release, how they can propagate, as well as possible means of intervention for combating them. Therefore, in general terms, the level of physical function provides the information that is necessary to identify accident mechanisms, to predict courses of events, and to judge the effects of countermeasures.

At the third level, general function, the accident potential is described in general terms that are relatively independent of the mechanism that initiated the incident. Thus, whether the accident occurred in one reactor subsystem or another is not of concern here. Rather, the accident is classified according to its potential for harming the public. For the radiological domain, we would like to know the amounts of direct exposure, ingested radiation, surface contamination, as well as the characteristics of the radiation release (gamma, radioiodine, etc.).

At the level of abstract function, we represent the information that is needed to compare, prioritize and coordinate the functions at the level below. Typically,

Table 1. The work of the center described as an abstraction hierarchy of risks and resources

	Domain of potential risk	Domain of mitigation resources
Functional purpose	Goal: To protect the citizens of the State in the event of a radiological incident Constraints: The resources and abilities associated with the organization	
Abstract function	Information used to compare, prioritize and coordinate lower level general functions Rules and regulations derived from laws dictated by central organizations	Criteria and measures for setting priority among resources (i.e. organizational units and/or services)
General function	Accident potential in general terms (e.g. direct exposure, surface contamination, different types of radiation, etc.)	Operational response functions: Informing public; Law enforcement Evacuation; Fire and rescue Take shelter; Emergency medical Food, water and milk; Social services Traffic/access control; Re-entry
Physical function	Different ways in which weather can affect the consequences of a release Various accident mechanisms leading to a release, propagation paths, and responses to possible means of intervention	Functional capability of equipment and tools, and professional ability of personnel for each agency involved in operational response
Physical form	Population density around each site Geographical distribution of meteorological conditions Description and location of systems, subsystems and components of each reactor (e.g. blueprints)	Locations and inventories of the various rescue mitigation resources (e.g. personnel, cars, trucks, aircraft, etc.) Service availability, delays, access and distances from site for each of the above

this information will take the form of rule sets and regulations that are dictated by organizational policy. In our case, the regulations that are used to set priorities take the form of guidelines suggested by the government plan for emergency management.

Finally, at the level of functional purpose, we have the basic purposes and constraints on decision making. In our case, the goal can be stated simply: To protect the public in the event of an emergency. The issue of constraints is a bit more involved. According to Rasmussen *et al.*,[7] the decisions made at this level determine the amount of risk which is acceptable, thereby defining the amount of funding and the strategies and resources made available for the emergency management effort. To a great extent, the government has already implicitly defined the level of acceptable risk by developing the plan for emergency management.

What can we gather from this domain representation? First, it is evident that the higher levels represent the problem with much less detail than the lower levels. For example, at the top level, we can simply describe the state of affairs in terms of whether the domain objectives are being met or not, whereas at the bottom level, the description would involve a myriad of details including, among other things, comprehensive descriptions of the status of many of the affected reactor's subsystems and the prevailing meteorological conditions. Clearly, some types of decisions require the detailed information presented at the lowest levels of the hierarchy, whereas other decisions can be made at the higher levels, where the overall functional objectives are of greatest concern.

The hierarchical representation in Table 1 also allows us to trace how the accidents can propagate with time. Accidents begin at the lowest level (e.g. a pump malfunctions), but their effects, if not counteracted, will be progressively revealed at higher levels. For instance, a certain malfunction of the plant (physical form), if not acted upon, can lead to a certain level of radiation inside the containment (physical function), which if aggravated by other factors can have the potential for a certain amount and type of airborne radioactivity (general function). Should such a release be likely, it will then require urgent attention in terms of what needs to be done (abstract function). Finally, if not effectively dealt with, the accident mechanism can make its present felt at the level of functional purpose in the form of a violation of the domain goal. This illustrates how events can propagate in a bottom-up manner through the hierarchy if they are not effectively contained.

Typically, decision making activities will cross the different levels of the hierarchy. For example, the level of general function can serve the level of abstract function by providing a description of the projected state of the incident in functional terms (i.e. independent of the specific event which initiated the incident). This status information will then serve as input to the guidelines which are the realization of the higher order goal of protecting public safety.

What practical implications can be derived from the domain representation illustrated in Table 1? First, the domain representation distinguishes different aspects of the work domain. This allows us to identify the demands that are being made on personnel. Secondly, the representation allows us to determine what information is needed at the various levels of the hierarchy. This information serves as the input to the decisions that need to be made throughout the hierarchy, and therefore specifies the informational content of the computer database that needs to be available to REAC personnel. Information needs can vary considerably as a function of where one is located in the space shown in Table 1. Thirdly, the domain representation provides a basis for organizing the database of information. Because we have a description of the functional structure of the domain, we can group items of information that belong together (i.e. that refer to the same level or function).

In order to have any practical impact, however, information needs to be translated into action. An appropriately designed computer system can facilitate the mapping between incoming information and available options for action. The other dimensions of the model mentioned above (i.e. the decision ladder and organizational analyses) together with the means-end hierarchy representation provide a strong basis for developing computer support systems that satisfy this objective.

To summarize, we have found that the conceptual tools developed by Rasmussen are very useful in analyzing the deep structure of a complex work domain. In particular, this set of tools allows us to develop a general model that is independent of the particular details of a given scenario and the subjective work preferences of a given set of personnel. An effective control room design should be flexible enough to support any well-trained set of operators in coping with the contingencies of any given scenario. This implies that a cognitive task analysis should be general enough to capture this wide range of demands, instead of being based on some 'normative' work procedure and a small number of predictable incidents. The general model briefly described above is directed at satisfying these objectives.

By the end of the project, after 18 months of interacting with the personnel at the Center, we concluded that the generic model is very useful for coordinating the thinking of the analysts. It is also very useful in planning improvements in the design of displays. For example, the results of a task analysis of the environmental analysts's and radio dispatcher's jobs, put into the context of a hierarchical means–ends analysis, provided a strong organizing principle for the redesign of displays. The recommendations were adopted and appeared to be successful. On the other hand the model played no significant role in the more general recommendations about changing the layout of the room, or in analysing problems of communication with outside agencies. If conceptual frameworks such as the Rasmussen ladder are to be of practical use to designers, much work remains to be done to discover to which sorts of problems they are most applicable, and exactly how they can be used in practical design decisions. We do not yet understand the scope of their domains.

VALIDATING THE TASK ANALYSIS

If we combine the four approaches which we have described, we believe that we have the basis for a fairly complete cognitive task analysis of a team of operators. The remaining question is: How to validate the cognitive task analysis?

We will be adopting two complementary methods. One approach is to validate via prediction. The analysis which we have conducted clearly indicates certain decision points which are particularly sensitive to certain sources of information, the integrity of certain communication channels, and patterns of interaction within the Center. Given that exercises are conducted several times a year, it should be possible to predict what particular problems in equipment, communication, information acquisition, etc., will be most critical for the operation of the Center. One or two of these could be embedded in the exercises on each occasion, and over a period of a year or two a substantial number of the predictions tested. A virtue of this approach is that it allows a concurrent evolutionary development of the facility, so that where predictions are correct, steps can be taken to improve the design and operations of the Center.

Another approach to validation is to use the general model to interpret and describe a number of exercises. If effective, the model should provide an effective 'language' for describing the operations that are observed under a wide variety of conditions. To the extent that is so, and to the extent that the observed patterns of work are consistent with the descriptions provided by the general model, the analysis can be said to be validated.

CONCLUSIONS

Cognitive task analysis of team work activities is extremely difficult, time consuming and labor intensive. We believe it to be possible along the lines we suggest. We further believe that every effort must be made to develop and improve on these and related techniques, for the chief characteristic of advanced technological systems is that traditional behavioral task analysis at the level of an individual is wholly inadequate to allow safe and rational design and operation of human–machine systems.

ACKNOWLEDGEMENTS

This work was supported by the Illinois Department of Nuclear Safety. Our thanks are due to all personnel for their assistance and cooperation.

REFERENCES

1. Rasmussen, J., *Information Processing and Human–Machine Interaction: An Approach to Cognitive Engineering*. North-Holland, New York, 1986.
2. Woods, D. D. & Hollnagel, E., Mapping cognitive demands in complex problem-solving worlds. *International Journal of Man–Machine Studies*, **26** (1987) 257–75.
3. De Keyser, V., Decortis, F., Housiaux, A. & Van Daele, A., *Les communications hommes–machines dans les systémes complexes: Caractéristiques et utilisation des informations centralisees*. Université de Liège, Service de Psychologie du Travail, Liège, Belgium, 1987.
4. Rochlin, G. I., La Porte, T. R. & Roberts, K. H., The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review* (Autumn 1987) 76–90.
5. Meister, D., *Behavioral Analysis and Measurement Methods*. John Wiley, New York, 1985.
6. Rasmussen, J., Concepts for analysis of cooperative work. Unpublished working paper, Risø National Laboratory, 1989.
7. Rasmussen, J., Pedersen, O. M. & Grønberg, C. D., *Evaluation of the use of advanced information technology (expert systems) for data base system development and emergency management in non-nuclear industries*. Risø-M-2639, Risø National Laboratory, Electronics Department, Roskilde, Denmark, 1987.