DTU Library

The role of hierarchical knowledge representation in decisionmaking and system management

Rasmussen, Jens

Published in:
IEEE Transactions on Systems, Man and Cybernetics

Link to article, DOI:
10.1109/TSMC.1985.6313353

Publication date:
1985

Document Version
Peer reviewed version

The Role of Hierarchical Knowledge Representation in Decision making and System Management

JENS RASMUSSEN, SENIOR MEMBER, IEEE

Abstract—The knowledge representation of a decisionmaker in control of a complex system can be structured in several levels of abstraction in a functional hierarchy. The role of such an abstraction hierarchy in supervisory systems control is reviewed, and the difference between causal and intentional systems and formal games in terms of the role of an abstraction hierarchy in the related decision strategies are discussed. This relationship is then discussed with reference to the classical psychological problem-solving research of Selz el al. Finally, the implications for design of decision support systems are discussed. It is argued that an explicit description of the functional properties of the system to be controlled in terms of an abstraction hierarchy is necessary for a consistent design of data bases and display formats for decision support systems. Also, it is necessary to consider the role of the abstraction hierarchy in reasoning when planning experiments on human decision making.

I. INTRODUCTION

DEClSlONMAKlNG in supervisory control of complex systems is typically related to singular situations with major disturbances, technical faults, inappropriate human actions, and combinations of such events. For such situations the established control algorithms do not apply, and planning of proper control actions by the controller depends on knowledge about the functional properties of the system for ad hoc diagnosis, evaluation, and planning. The decisionmaker has to adapt to the requirements of the system under the specific abnormal conditions. Consequently, the information processes that will be applied to develop the proper control actions depend on characteristics of the specific situation and on the individual decisionmaker.

This means that it is difficult, if not impossible, to develop models that make it possible to predict the information processes during a particular rare event and that can be used as a basis for the design of decision support and interface systems. Rather, the design will be in terms of a functional envelope, which leaves room for the individual decisionmaker to adapt to situations within the categories adopted as design basis. For such an envelope design, a model of the decision process itself is not necessary; instead design can be based on higher-level structural representations of major characteristics separately, such as:

- the control requirements of the system for relevant categories of situations;
- the decision context or problem space, i.e., a systematic representation of the functional properties of the system;
- a repertoire of possible and effective strategies for the various phases of decision making, such as diagnosis, evaluation, and planning; and
- a representation of information processing capabilities and limitations of the decisionmaker and of the subjective formulation of goals and criteria for choice among possible strategies, i.e., the human product and process criteria.

In the present discussion focus will be on the properties of a hierarchical representation of the problem space.

II. HIERARCHICAL KNOWLEDGE REPRESENTATION

Control of complex systems depends on the means for coping with the complexity. However, from the controller's point of view, the complexity of a system is not an objective feature of the system [1]. The complexity observed depends upon the resolution applied for information search. A simple object becomes complex if observed through a microscope. Objective complexity can only be defined with reference to a given representation of a system. Therefore the complexity faced by controllers is determined by the representation of the internal state of the system, which the interface allows the controller to develop for the various work conditions. This means that the apparent complexity of a system ultimately depends on the technology of the interface system. For instance, the complexity of the traditional industrial control consoles depends on the one-sensor/one-indicator technology. Only one level of resolution of the representation is available, and this has to meet the most detailed level needed in any situation. In that case the interface must be complex by the law of requisite variety. The only way to cope with control of systems that are complex in terms of large numbers of information sources and devices for basic control actions is to structure the situation and thereby transfer the problem to a level with less resolution.

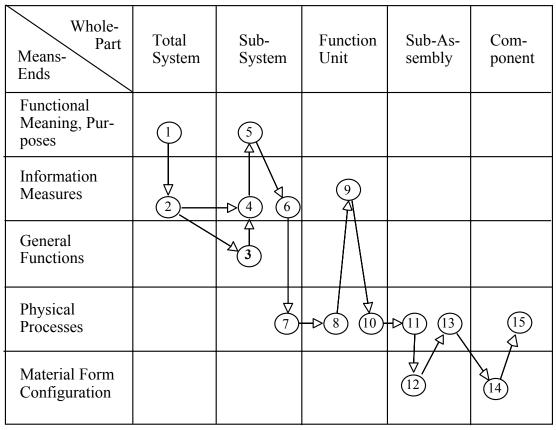

Fig. 1. The focus of a computer trouble shooter's attention can vary with respect to the part/whole dimension and to the level of abstraction. The trace followed in an actual diagnostic task is illustrated.

From analysis of verbal protocols recorded by people working on complex systems (process plant operators and computer maintainers [2], [3]), we have found that this structuring can typically be done along two dimensions: one is the whole/part consideration, in which the system can be seen as a number of related components at several levels of physical aggregation; another is the level of abstraction in the representation, i.e., the degree to which the physical implementation of functions is maintained in the representation. Very often, change of the level of physical decomposition within the span of attention is coupled with a change in the level of abstraction in representation, but, basically, the whole/part and the abstract/concrete dimensions are conceptually separate. As an example, the changes in the two dimensions during trouble shooting in a computer system are shown in Fig. 1.

In the following, we will mainly consider the role of the abstract/concrete dimension in knowledge representation.

III. THE ABSTRACTION HIERARCHY

In the abstraction hierarchy, as shown in Fig. 2, the functional properties of a technical system are represented in several levels of functional abstraction along the means-end dimension.

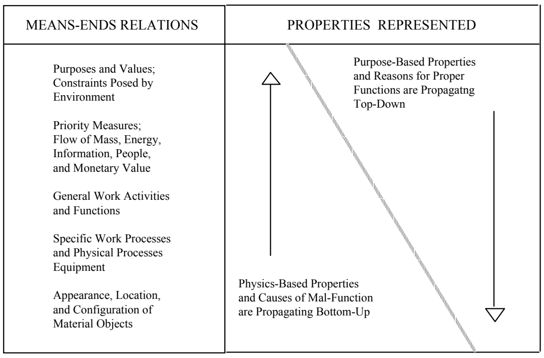

Fig. 2. The system properties considered can be described at various levels of abstraction, representing the physical implementation and functional purpose in varying degrees.

The number of levels between the ultimate purpose of the system and the material physical implementation which will be relevant depends on the type of system and the aim of the study. The levels shown in Fig. 2 have been identified from verbal protocols related to energy-conversion systems and digital computers, which have in common the feature that the functions can be represented in purely abstract, i.e., process- and equipment-independent terms, related to energy and information flow topology, respectively. However, it is typical that the levels of Fig. 2 all refer to specific categories of engineering languages for system description [3]-[5].

The lowest level of abstraction represents only the physical form of the system: its material configuration. The next higher level represents the physical processes or functions of the various components and systems in a language related to their specific electrical, chemical, or mechanical properties. Above this, the functional properties are represented in more general concepts without reference to the physical process or equipment by which the functions are implemented, and so forth. At the lower levels, elements in the process description match the component configuration of the physical implementation. When moving from one level of abstraction to the next higher level, the change in system properties represented is not merely the removal of details of information on the physical or material properties. More fundamentally, information is added about higher level principles governing the cofunction of the various functions or elements at the lower level. In man made systems these higher level principles are naturally derived from the purpose of the system, i.e., from the reasons for the configurations at the level considered. Change of level of abstraction involves a shift in concepts and structure for representation as well as a change in the information suitable to characterize the state of the function or operation at the various levels of abstraction. Thus a supervisory controller will ask different questions about the state of a system, depending on the task and the nature of the currently active internal representation.

In other words, models at low levels of abstraction are related to a specific physical world that can serve several purposes. Models at higher levels of abstraction are closely related to a specific purpose that can be met by several physical arrangements. Therefore shifts in the level of abstraction can be used to find suitable paths for the transfer of knowledge from previously experienced cases and problems. At the two extreme levels of models the directions of the paths available for transfer are in a way orthogonal, since transfer at the lowest level follows physical material properties, while at the highest it follows purpose.

In the present context an important use of the abstraction hierarchy is as a framework for describing the control tasks required to maintain satisfactory system operation. States can only be defined as errors or faults with reference to the intended functional purpose. Causes of improper functions depend upon changes in the physical or material world. Thus they are explained "bottom-up" in the levels of abstraction. In contrast, reasons for proper function are derived "top-down" from the functional purpose. This distinction is illustrated in Fig. 2.

During system operation the task of the control system (which includes the automatic control system as well as human operators) will be to ensure, by proper actions on the system, that the actual state of the system matches the target state specified by the intended mode of operation. This task can be formulated at any of the different levels of abstraction. During system start-up, for instance, the task moves bottom-up through the hierarchy. In order to have an orderly synthesis of the overall system function during start-up, it is necessary to establish a number of autonomous functional units at one level before they can be connected to another unit at the next higher level. This definition of autonomous functional units at several levels is likewise important for the orderly breakdown of system functions for shutdown and for reconfiguration during periods of malfunction. For such considerations there will typically be a tight coupling between the means / end (concrete / abstract) and the part / whole dimension.

During emergencies and major disturbances an important control decision is to prioritize by selecting the level of abstraction at which the task should be initially considered. In general, the highest priority will be related to the highest level of abstraction. First, judge the overall consequences of the disturbance for the system function and safety in order to see whether the mode of operation should be maintained or switched to a safer state (e.g., stand-by or emergency shutdown). Next, consider whether the situation can be counteracted by reconfiguration to use alternative functions and resources. This is a judgement at a lower level of functions and equipment. Finally, the root cause of the disturbance is sought in order to determine how it can be corrected. This involves a search at the level of physical functioning of parts and components.

When a disturbance has been identified and the control task located at a level of abstraction depending upon the situation, the supervisory control task includes the determination of the target state derived top-down for the chosen operating mode. In addition, the available resources for reconfiguration and limits of capabilities must be derived from levels below. A decision task in a particular situation can be formulated, as has been discussed, with reference to a process at any level of the abstraction hierarchy. For the task the operator will typically need information from the level considered, representing the functional structure and state of this process, i.e., regarding what is controlled. But he will also need information from the level above, which is related to the immediate purpose of the control decision, i.e., why it is made, as well as from the level below, i.e., how a decision can be implemented [1].

IV. CAUSAL, INTENTIONAL, AND FORMAL SYSTEMS

The line of reasoning in the above discussion reflects the engineers' reas-oning about man-made systems: the reason for the functional structure of a system is derived top-down from the purpose, and the response of the system to an external influence is predicted bottom-up from causal laws. Such systems can be called causal systems.

Not all systems, however, can be modeled this way. Systems with a high degree of autonomous internal functioning, with self-organizing and highly adaptive features, may change their internal functional organization frequently in order to meet the requirements of the environment and to suit their internal goals or performance criteria. Even though such systems are basically causal and controlled by laws of nature, their complexity makes it impractical, if not impossible, to explain or predict their performance by functional analysis during real-life decision making. The alternative is to consider such systems as intentional systems, controlled in their response to external influence within their range of capability by their "intention" to act derived from the individual value structure and internal goals. Prominent examples are humans and social systems. For humans, an abstraction hierarchy similar to that of Fig. 2 is readily suggested with the different levels representing models in terms of anatomy, physiology, psychology, information processing, and value structures such as myths and religions [4]. A classical discussion since Stuart Mill is whether the problem of relating mental events to physiological states is due only to the complexity of the human organism and will be resolved by means of natural science in terms of causal explanations bottom-up in the hierarchy, or whether, in principle, basically different intentional models are needed [6].

It is, however, not only a problem in living organisms. Many technical systems such as control systems and information-processing systems are very complex and have no simple relationship between their basic physical processes and their function in the information domain. Therefore predictions regarding their behavior are more readily made when considering the systems as intentional systems [7]. Even in the case of relatively simple systems, operators can be seen in verbal protocols to develop an explanation of system behavior from a top-down "redesign" of a reasonable functional structure from its supposed purpose, rather than to collect information on its actual physical structure.

Decision making in control of intentional systems is based on knowledge about the value structures of the system, the actual input from the system's environment, and its internal limiting properties; i.e., it is based on reasoning topdown in the abstraction hierarchy with little or no consideration of the internal causal structures or functions. This is probably the reason why top-level executive decisionmakers, according to Mintzberg's study [8], do not behave according to analytical decision models, but prefer live action and social contacts to the analysis of abstract information and current data— even gossip and hearsay— instead of statistics and status reports. Meeting people and considering hearsay are probably the best sources of information on current trends in value structures. In addition to this need for top-down reasoning in managerial decision making, the reluctance to follow the rational normative decision models is related to the replacement of a novice's rational strategies by intuitive judgements based on situational and concrete evidence as expertise is developed. This has been emphasized by Dreyfus [9].

The strong emphasis on causal relationships in the modeling of social systems, such as Forrester's world model [10], rather than careful consideration of the dynamics of value structures may greatly decrease the quality of long-term projections. A similar critique of traditional historical theories has been expressed by Toynbee [11], who has a system-oriented approach to the history of societies. He defines a society in the following way [11, p. 43]:

A society is the total network of relations between human beings. The components of society are thus not human beings but relations between them.

Presenting his "challenge-and-response model," he argues 1, p. 97]:

> In my search up to the present point, I have been experimenting with the play of soulless forces— vis inertia and race and environment—and I have been thinking in the deterministic terms of cause-and-effect. . . The effect of a cause is inevitable, invariable and predictable. But the initiative that is taken by one or the other of the live parties to an encounter is not a cause, it is a challenge. Its consequence is not an effect, it is a response. Challenge-and-response resembles cause-and-effect only in standing for a sequence of events.

The distinction between physical causal systems and intentional self-organizing systems must also be considered if results from research in human performance in games (for instance from artificial intelligence research) are considered for use in models of human interaction with physical systems. In two-person games like chess, a person faces a system that is not controlled by basic invariant laws but by the intentions and value structures of the opponent. The game itself only represents a means of communication, and the rules of the game serve only to constrain the decisions of the players to a well-defined set in each situation. Decisions depend upon prediction of the opponent's value structures and performance criteria and the strategy he adopts for the game. The difference between games like chess and other social-system contexts for management decision making is largely a question of formal consistency and invariance of the rule set. In games, the set of rules at the problem level is small and closed, and only the strategies for generation of proper action sequences are flexible and depend on top-down inferences regarding the opponent's intentions. In management decision making, there is room for the invention of new rules of the game within the constraints of legal rules and "fair play."

Formal systems for decision making are problems that are only defined at one single level of abstraction, such as geometrical theorem proving and construction, cryptographic problems, puzzles, and purely logic problems; see for instance Newell and Simon [12]. The problem is stated here as an initial state and a target state, and the task is to identify a sequence of allowed formal transformations which will close the gap. The category of formal systems also includes several of the "context-free" tasks, which are used for problem-solving and man-machine interface experiments. It should, however, be realized that problem-solving behavior may be very different in one-level formal systems and in a problem context of an abstraction hierarchy. This leads to a discussion of problem-solving strategies with reference to the means/end abstraction hierarchy.

V. PROBLEM-SOLVING STRATEGIES

An illustrative example of the role of the abstraction hierarchy can be found when comparing a decision task that has to be performed in a one-level formal description with the performance when the context is also available. The difference may partly be due to the use of shifts in the level of abstraction to find paths for transfer of solutions and strategies by analogy, but also due to the support of memory and search for rules in terms of structures at other levels of abstraction. A good empirical piece of evidence is the experiment made by D'Andrade (discussed by Rumelhart and Norman [13]) who repeats the experiment of Wason and Johnson-4aird [14]. This experiment was based on a set of cards which were hypothesized to represent the concept: if one side shows a vowel, then the back side displays an odd number. A subject was given a sample of four cards and asked which should be turned over in order to test the hypothesis. The experiment was repeated with the same concept disguised in a bill-signing context: if the amount of a bill exceeds $50, the supervisor must sign the back side. The ratio of correct solutions in the two experiments was 13 percent to 70 percent. Rumelhart and Norman conclude:

> What is the difference here? Why do people appear not to understand the meaning of it in the first case and understand it nearly perfectly in the second? This is exactly the kind of effect expected if our knowledge is embedded in a relatively inaccessible procedural format rather than as general rules of inference. The first case of the label factory represents a relatively unfamiliar case in which we cannot rely on specific knowledge and must, therefore, rely on general reasoning processes. The second case more nearly approximates our real-life" problem-solving situations. Once we can "understand" the situation, the conceptual constraints of our specific knowledge can be brought into play and the problem readily

> solved. It is as if our knowledge representation already contains all the reasoning mechanisms ordinarily required. Thus it would appear that the context dependencies inherent in the more procedural representational systems are also present in the human reasoning system.

It seems to be possible to translate this explanation directly into the terms of the present conceptual framework: in the first experiment the problem-solving is based on formal logical arguments at only one level of abstraction; on syllogistic logic, which requires manipulation of abstract symbols; and on storage of intermediate results in short term memory. In the second experiment the context defines an intentional system, in which the effects of the different decisions can be inferred very easily at the higher levels. The reasons for proper states can be inferred " top-down." The problem is solved by top-down model modification; that is, by transferring to a model of "reasonable states of affairs." Rumelhart and Norman refer to " understanding," which in our context can be viewed as the ability to transfer the problem—upwards towards a reason or downwards towards a cause—to a level where immediate intuition from experience is available.

The role of the abstraction hierarchy not only seems to be the transformation of a problem to a level at which a solution is more readily available, e.g., by use of analogical reasoning. The transformation between levels also seems to be a powerful tool for functional reasoning. Formal logical reasoning appears to consist of " horizontal" considerations within a single level of the hierarchy, based on the classical syllogistic reasoning about membership of exclusive categories. Practical functional reasoning is, on the other hand, related to "vertical" transformations in the means/end abstraction hierarchy.

The difference between classical formal logic and practical functional reasoning has long been a topic of discussion in philosophy; and, mentioned in passing, the well-structured nature of the abstraction hierarchy related to man made physical systems might make it a better-suited vehicle for resolving such questions than the all-encompassing " real world" normally considered by philosophers. Bosanquet [15] criticizes the classical syllogistic model of reasoning and, typically, selects as examples for his arguments electrical circuits with fuses and Harvey's discovery of the function of the circulation of blood. Bosanquet discusses at length the difference between syllogistic reasoning, which is linear, and functional inference, which he calls "systemic." He argues [15], [16] that explanation depends on the relational network of a whole, and that linear syllogistic arguments therefore are inadequate. Similar arguments have been presented recently by Harman [17]. He distinguishes inference or reasoning from proof or argument. To him, reasoning is a process of trying to improve one's overall view and is a holistic process that puts the problem into context, whereas rules of argument or proof are local rules of logical implication. He relates the distinction to a syntactic or grammatical basis, to a distinction between elements of logical form and nonlogical content—a distinction that is similar to the present discussion of arguments within a model at a single level of abstraction—and reasoning across levels by transformation and modification of models. Harman [18] argues that practical reasoning is concerned with what to intend and formal reasoning with what to believe. Formal logic arguments are *a priori* true or false with reference to an explicitly defined model, where functional reasoning deals with relationships between models, and truth depends on correspondence with the state of affairs in the real world.

Certain schools of psychology also emphasize the significance of the structure of a problem as a guide to thinking. Wertheimer [19] argues from an analysis of a number of problem-solving scenarios that problem-solving is a productive process which depends on the structure of the problem situation in a way which was neither considered in the then prevailing associative theories nor in classical logic. The process is not characterized by merely linking together elements piecemeal by associations. The flow is controlled by whole-characteristics, a "sensible expectation about structural truth." He argues the need for a theory that goes directly to the structural nature of the process and states " the gist of the thesis: structural reasons become the causes of the process," and he goes on to discuss the difference between reasons and causes in control of the process.

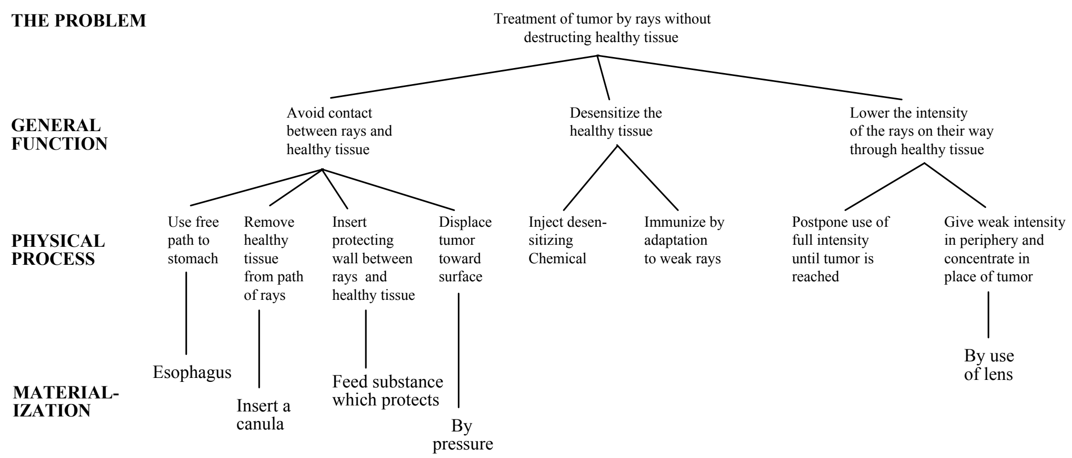

Figure 3. Generic solution tree found by Duncker in a problem solving task. Levels are similar to the abstraction hierarchy of figure 2. Adopted from Duncker [22].

Selz [20] (for a good English review see [21]) also argues against an association model of thinking. His work is interesting in the present context, since his arguments are based on an analysis of errors in thinking. His premise is that, given that the mechanism is associative, errors should include a large fraction of associations with no functional meaning or without connection to the task. This is not what he found. On the contrary, errors typically appeared to be results of solution trials with regard to the task that is somewhat misconceived. Selz is very modern in his conception of problem-solving procedures which are determined by: l) the intellectual personality, i.e., the repertoire of solving operation dispositions; 2) the features of the problem; and 3) the subject's intention. The course of a thought process is controlled by the subject's "schematic anticipation," involving gap, tension, and four basic operations, such as likeness evocation, abstraction, combination, and complex completion, which can all be readily related to operations in an abstraction hierarchy. Selz's distinction between " productive" and " reproductive" thinking is related to knowledge/rule-based control of behavior [4]. In productive thinking he distinguishes between "finding the means" and "applying the means." Finding new means may involve: l) reproductive abstraction of means, which identifies the means by top-down search in the hierarchy (compare with Duncker's "by means of which" relation, mentioned subsequently); 2) coincidental identification, which looks like Duncker's "suggestion from below," and 3) identification of means from structural insight into the nature of the task, i e., restructuring through understanding. In all, the work of Selz points to the importance of a well-structured representation of the problem context. Through de Groot's [21] study of chess strategies, the work of Selz has influenced modern artificial intelligence (Al) research, but mainly through work on games and formal problems [12].

The role of a multilevel abstraction hierarchy in problem-solving is most explicitly seen in Duncker's [22] research on practical problem-solving related to physical causal systems (radioactive tumor treatment and temperature compensated pendula). Based on verbal protocols, Duncker describes how subjects go from the problem to a solution by a sequence of consideration, where the items proposed can be characterized by a " functional value" feature pointing upwards to the problem, and a " by means of which" feature pointing downwards to the implementation of a solution (see Fig. 3). The relation to the abstraction hierarchy as shown in Fig. 2 is clear. He states:

> The final form of a solution is typically attained by way of mediating phases, of which each one in retrospect possesses the character of a solution, in prospect, that of a problem...when one closely examines what Maier calls direction (in thinking), it becomes clear that direction is nothing but the earliest phase of the solution, i.e., the reformulation of the problem as it initiates the solution process concerned.

This means that direction in thinking is given by the structure of the available means-end abstraction hierarchy with its gaps; compare this with the structural anticipation of Selz.

Yet another observation on the role of an abstraction hierarchy on understanding a mechanical device has been reported by Rubin [23], who reports an analysis of his own efforts to understand the function of a mechanical shutter of a camera. He finds that consideration of purpose or reason plays a major role in the course of arguments. He conceived all the elements of the shutter in the light of their function in the whole. He did not perceive the task to explain how the individual parts worked, but rather what their functions were in the whole. How they worked was immediately clear when their function was known. He mentions that he finds it an analytical task to identify the function of parts, the direction of thought being from overall purpose to the individual function (top-down considerations). The hypothesis necessary to control the direction is then readily available. This approach was found to have additional advantages: solutions of subproblems immediately have their place in the whole picture, and it is immediately possible to judge whether a solution is correct or not. In contrast, arguing from the parts to the " way they work" is much more difficult, this being a synthesis. Solutions of subproblems must be remembered in isolation, and their correctness is not immediately apparent.

For comparison, readers with philosophical inclinations should read Kant's discussion of the problem faced by "dissectors of plants and animals" [24]:

> In fact, they can as little free themselves from this teleological proposition as from the universal physical proposition; for as without the latter we should have no experience at all, so without the former we should have no guiding thread for the observation of a species of natural things which we have thought teleologically under the concept of natural purposes.

VI. IMPLICATIONS FOR THE DESIGN OF DECISION SUPPORT SYSTEMS

The task of a supervisory controller during new and unfamiliar system disturbances is very similar to the problem-solving task studied by, for instance, Duncker: the problem is formulated during the diagnostic phase, and a solution is sought in terms of a reconfiguration of the available physical resources. The preceding discussion of the importance of the structure of the problem in terms of a means-end hierarchy therefore has a number of implications for analysis and design of systems for support of supervisory systems control.

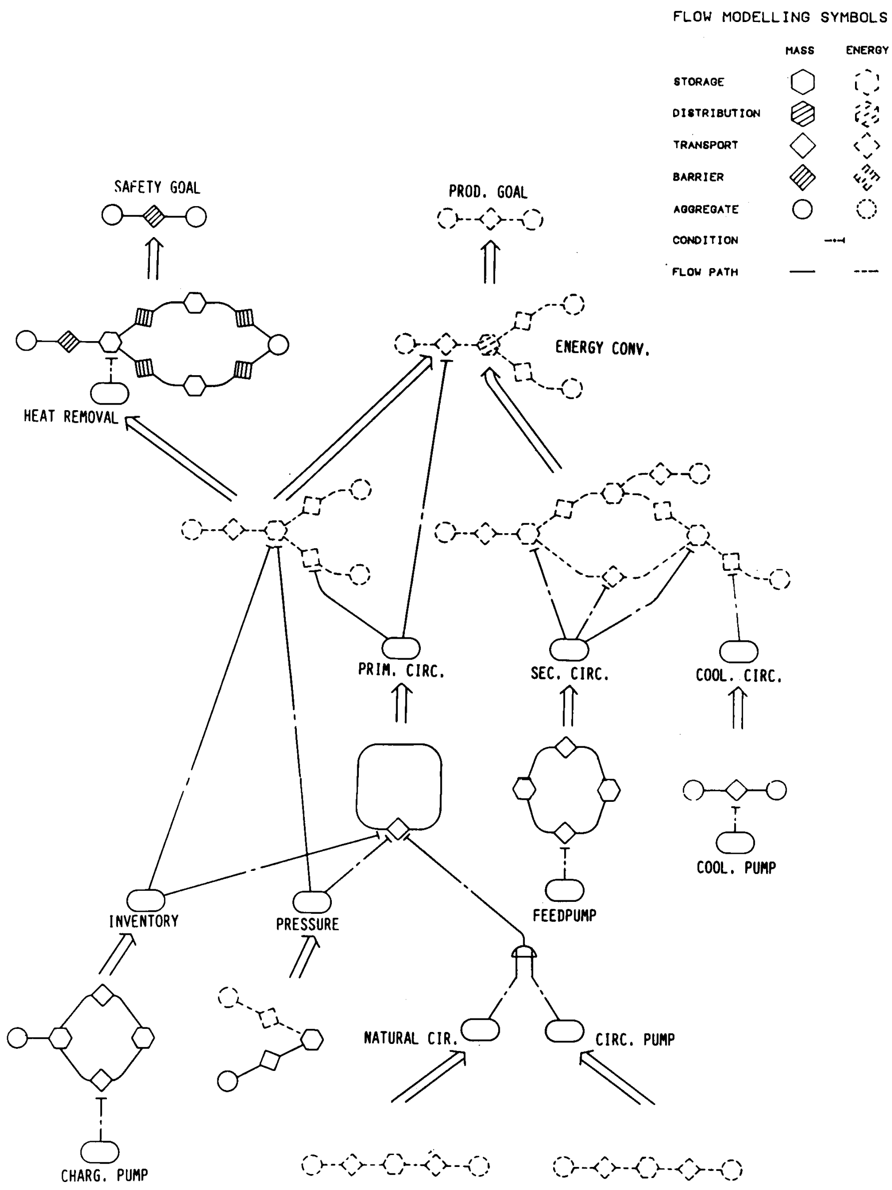

Figure 4. Illustration of many-to-many mapping between levels of purpose/functions/equipment of a nuclear power plant and the use by Lind [25] of mass/energy flow topology for a systematic representation of the relationships.

Structuring Representation in Data Bases: In the design of man-machine interfaces for supervisory control as, for instance, industrial process control consoles, the emphasis has traditionally been on the presentation of measured data representing the physical state of the system and its processes. In complex systems this information has been supplemented with information about the underlying functional structure by graphical means such as mimic diagrams, etc. This information is intended to serve as the controller's identification of the actual state of processes "bottom-up" through the hierarchy. Infor-mation representing the intentions behind the system design, in terms of the purpose of functions and equipment and of constraints upon the acceptable operation, as derived from safety considerations, for instance, has only been very sparsely represented. This means that the interface system gives little or no support in the "top-down" derivation of the proper or acceptable states of processes. This kind of information was supposed to be immediately available to operators from their basic training. However, as systems become complex and potentially risky, and very low probabilities of erroneous deci-sions during rare events are required, this can no longer be assumed. During such sit-uations information about the purpose of interlocks, properties of equipment if used for untraditional purposes, etc., may be vital for *ad hoc* improvizations. Consequently, it becomes increasingly important to include in the support for decision making the information needed for "top-down" consideration of reasons. This means that the properties of the system to be controlled should be represented in terms of a consistent means/end hierarchy, systematically mapping the purpose / function / equip-ment relationships.

This leads to two problems: to find the information needed and to structure the representation. Information on " reasons" behind design decisions is to a great extent implied in standard practice, or only their implications are recorded by the designer in specifications and drawings. Regeneration may require quite a fair amount of work. For the structuring of the description, a systematic method based on a formal lan-guage is required. An approach to this problem has been taken by Lind [25], who has developed a multilevel description of process systems in terms of their mass and energy flow topology (see Fig. 4). This de-scription is basically a representation of the flow topology of the system at several levels of decomposition, but it maps very well onto a means/end hierarchy for a given system and is well suited for structuring a data base for supervisory support systems independently of the formats chosen for information presentation. In particular, this formal representation supports systematic analysis for the information gathering from the designers, and, in addition, it will serve as a basis for automatic inference generation in the process computer's processing of the plant information.

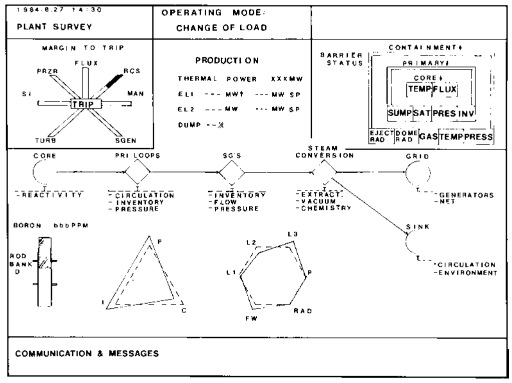

(a)

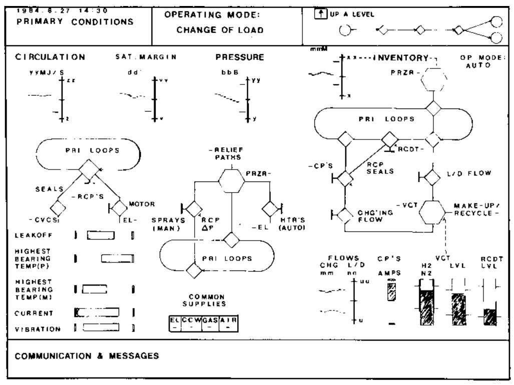

(b)

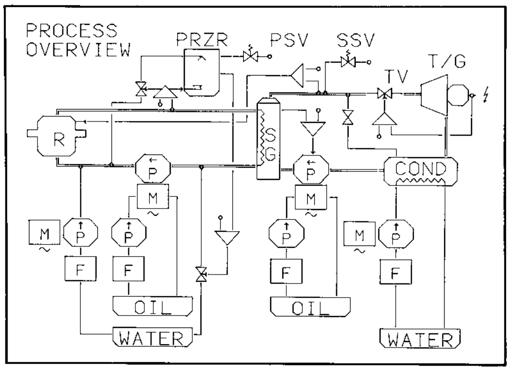

(c)

Figure 5. Computer-generated displays can be matchted with the content and form of useful mental models at the different levels of abstraction. Display formats for control of a simulated nuclear power plant, designed by Goodstein [27] for experimental studies, are illustrated. (a) Overall energy flow of the plant. (b) Lower level functions. (c) Basic physical anatomy as a format for displaying sets of equipment status data.

Structuring Information Presentation: Given that the information is present in the data base in a structured way, the problem is how to present it. During the various phases of a supervisory control task, the relevant level of consideration will vary in the means/end hierarchy. An obvious proposal will be that the data from the system should be available in preprocessed form, matching the level considered. For this purpose the formal flow term representation is well suited for data integration by the computer. In addition, the display format should match a useful mental model of the structure of the functions at that level. In this way the information processing required by a supervisory controller for preparation of data to match his decision task can be greatly reduced. However, there is a price, which is the added information retrieval task of finding the relevant display in perhaps a large library. An experimental program is needed in order to analyze problem-solving behavior and subjective preferences in a task environment where information is available at several levels of a means/end hierarchy and in several formats of presentation. Goodstein *et al.* [26] have developed an experimental set-up based on a simulated nuclear power plant for this purpose. For these experiments a set of displays based on the multilevel flow concept has been designed [27] (see Fig. 5).

Predicting Decision Error Modes: Another important implication of the functions of the means/end hierarchy in the control of problem-solving behavior is its influence on error modes during decision making. As Selz noticed, decision errors are not stochastic events but depend upon the structure of the problem space in question. If supervisory decision making is considered as resource management in a task performed in a purpose/function/equipment hierarchy, several kinds of interferences can be suggested as sources of systematically appearing decision errors. Equipment may be useful for different purposes in different situations or during different tasks, and conflicts may appear between the efforts of different users or the aims of the same person at different times. Seemingly unexplicable human acts during a critical task may be caused by mistakes caused by similarities of features at one or another level in representation. Similarly, decisions may be judged erroneous in retrospect, but in the actual situation may be caused by attempts to test a very reasonable but wrong hypothesis.

The conclusion of this is that a systematic representation of the means/end relationships of the control object of a supervisory decisionmaker is a necessary prerequisite for modeling and predicting decision errors.

Planning of Experiments: The distinction between problem-solving which is performed by formal logic within a problem space of only one level of abstraction and problem-solving based on transformation through the levels of an abstraction hierarchy has immediate implications for laboratory experiments on decision making. Frequently, more or less context-free problem representations are used for experiments on problem-solving strategies. These experiments are very effective when the aim is to model logical reasoning in a closed formal system; but the implications for real-life tasks may be difficult to establish, since the setting does not invite a shift in the level of abstraction and hence does not support problem-solving from prior experience based on functional inference and analogies. Furthermore, the lack of context leads to the replacement of the task to infer proper states from high-level purposes and the effect of complex disturbances from physical relationships by using the formal arguments of the one-level rules of the game. However if embedded in various context scenarios, a context-free system simulation can be a very flexible experimental tool. Adding context to an experiment may change the task in two ways. It may add to the task through the transformation of high-level criteria and physical disturbances into formal parameters, and it may change the decision strategy by supporting multilevel inferences as illustrated in the D'Andrade case. This means that, in addition to the traditional experimental psychology methodology, there is a need for experiments in much more complex and yet controlled settings [26].

CONCLUSION

The conclusion of this discussion is that the designers of industrial process systems have to analyze and formulate very explicitly the functional properties and control requirements of the systems. This is necessary for the design of decision support systems for supervisory control based on modern information technology. Descriptions in terms of an abstraction hierarchy representing the many-to-many mapping in the purpose / function / equipment relationships of complex systems seem to be a useful interface between technical control requirements and research on human problem-solving and, therefore, a useful tool for systems design.

REFERENCES

[l] J. Rasmussen and M. Lind, " Coping with complexity," presented at the European Ann. Conf. Human Decision Making and Manual Control, Delft, 1981. (Also in Risø Nat. Lab., Roskilde, Denmark, Rep. No. M-2293).

[2] J. Rasmussen and A. Jensen, "Mental procedures in real life tasks: A case study of electronic trouble shooting," *Ergonomics, vol.* 17, no. 3, pp. 293-307,1974.

[3] J. Rasmussen, "On the structure of knowledge—A morphology of mental models in a man-machine context," Risø Nat. Lab., Roskilde, Denmark, Rep. No. M-2192, 1979.

[4] J. Rasmussen, "Skills, rules, knowledge; signals, signs, and symbols; andother distinctions in human performance models," *IEEE Trans.Syst. Man, Cybern., vol.* SMC-13, no. 3,1983.

[5] J. Rasmussen, "Strategies for state identification and diagnosis," in *Advances In Man-Machine Systerns Research, vol. 1, W.* B. Rouse, Ed. Greenwich, CT: J.A.I. Press, 1984.

[6] P. Winch, *The Idea of a SoNal Science.* London: Routledge & Kegan Paul, 1958.

[7] D. C. Dennett, "Intentional systems," *J. Philos., vol.* LXVIII, no. 4, Feb. 25, 1971.

[8] H. Minzberg, *The Nature of Managerial Work.* New York: Harper and Row, 1973.

[9] S. E. Dreyfus and H. L. Dreufus, "A five-stage model of the mental activities involved in directed skill acquisition." Operations Res. Center, Univ. California, Berkeley, ORC-80-2, 1980.

[10] J. W. Forrester, *World Dynamics.* Cambridge, MA: Wnght-Allen, 1971 .

[11] A. Toynbee, *A Study of History.* Oxford: Oxford Univ. and Thames and Hudson, 1972.

[12] A. Newell and H. A. Simon, *Human Problem Solving.* New Jersey: Prentice-Hall, 1972.

[13] D. E. Rumelhart and D. A. Norman, "Analogical processes in learning," in *Cognitive Skills and Their Acquisition,* J. R. Anderson, Ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1981.

[14] P. C. Wason and P. N. Johnson-Laird, *Psychology of Reasoning.* Cambridge, MA: Harvard, 1972.

[15] B. Bosanquet, *Implication and Linear Inference.* London: MacMillan, 1920.

[16] B. Bosanquet, *Logic or the Morphology of Knowledge.* Oxford: Clarendon, 888.

[17] G. Harman, " Logic, reasoning, and logic form," in *Language, Mind, and Brain, T. W.* Simon and R. J. Scholes, Eds. Hillsdale NJ: Lawrence Erlbaum Associates, 1982.

[18] G. Harman, "Practical reasoning," *Reu. Metaphys., vol.* XXIX, no. 3, pp. 431-463, 1976.
[19] M. Wertheimer, *Productive Thinking.* New York: Harper and Brothers, 1945.
[20] 0. Selz, *Zur Psychologie des produktiven Denkens und des Irrtums.* Bonn: Friederich Cohen, 1922.
[21] A. D. de Groot, *Thought and Choice in Chess.* Hague: Mouton 1965.
[22] K. Duncker, " On problem solving," *Psychol. Monog., vol.* 58, no. 5,
[23] E. Rubin, "Vorteile der Zweckbetrachtung fur die Erkentnis," Z. *Psychol., vol.* 85, 1920. (Also in: *Expenmenta Psychologica.* Copenhagen: Munksgaard, 1949.)
[24] T. M. Greene, Ed., *Kant Selections.* New York: Charles Scribner's Sons, 1929.
[25] M. Lind, " Multilevel flow modelling of process plants for diagnosis and control," presented at the Int. Meeting Thermal Nuclear Power Safety, Chicago, IL, August 29-September 1,1982.
[26] L. P. Goodstein *et al.,* "The GNP testbed for operator support evaluation," Riso Nat. Lab., Roskilde, Denmark, Rep. No. M-2460 1983.
[27] L. P. Goodstein, "Information presentation for decision making," Riso Nat. Lab., Roskilde, Denmark, Rep. No. M-2461, 1984.