AUGUST 1987 LIDS-P-1688

EVALUATION OF DECISION AIDING IN SUBMARINE EMERGENCY DECISIONMAKING*

by

Scott T. Weingaertner **
Alexander H. Levis***

ABSTRACT

An analytical approach is presented for investigating the quantitative effects of a decision aid upon user workload and organization performance. The approach is applied to the problem of submarine emergency control. Results of the analysis indicate the importance of considering the organizational context when contemplating the design of decision aids and illustrate the implications for workload of *meta-decisions* about whether and how decision aid information is to be used.

* This work was carried out at the MIT Laboratory for Information and Decision Systems with support provided by the Charles Stark Draper Laboratory, Inc. under contract DL-K-260948 and the Joint Directors of Laboratories under contract N00014-K-85-0782 from the Office of Naval Research.

**ANSER, Inc., 1215 Jefferson Davis Hwy., Arlington, VA 22202.

***Laboratory for Information and Decision Systems, MIT, Cambridge, MA 02139

EVALUATION OF DECISION AIDING IN SUBMARINE EMERGENCY DECISIONMAKING

Scott T. Weingaertner

ANSER, Inc.
1215 Jefferson Davis Hwy
Arlington, VA 22202

Alexander H. Levis

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

An analytical approach is presented for investigating the quantitative effects of a decision aid upon user workload and organization performance. The approach is applied to the problem of submarine emergency control. Results of the analysis indicate the importance of considering the organizational context when contemplating the design of decision aids and illustrate the implications for workload of *meta-decisions* about whether and how decision aid information is to be used.

INTRODUCTION

With improvements in information processing technology has come an enormous amount of effort to design systems for helping individuals make decisions. Decision aids and support systems are often sought as a fix to existing overload and performance problems. The aim of this paper is to place the design of such systems in a larger context than perhaps is usual in order to reveal caveats for designers to consider.

Too narrow a focus in the design of decision aids may overlook two subtle but important issues: (i) decisionmakers often function as part of an organization whose output is supremal to that of the individual decisionmakers (DMs) comprising it, and (ii) the individual DMs may often be faced with *meta-decisions* about whether and how to use the information provided by an aid. The former issue implies that a decision aid should not be designed or put into use without considering its effects on the organization taken as a whole. The latter implies that the presence of the aid subtly changes the decision faced by the user, a special concern if the organization is an existing one to be retrofitted with an aid. This paper suggests an approach to the modeling and analysis of decision aids that not only allows their effects upon users to be characterized, but addresses the two issues identified above. An application of the approach to an actual system is presented wherein a five DM organization is modeled with and without a decision aid. The results generated underscore the importance of considering the issues of organization architecture and meta-decisions when developing decision aids.

The approach taken in this research builds upon the work of Levis, et al, in which an information theoretic model of the human DM permits the explicit and detailed representation of internal decision algorithms and strategies as well as external interactions between all the members of the organization, whether human or machine, and where all interactions, both internal and external, are represented using the Petri Net formalism [1], [2], [3], [4], [5], [6]. The measures generated by this methodology are individual DM workload, modeled as information processing activity, and organization performance, usually the expected value of the error. Application of this methodology yields an organization model that allows for the generation of both numerical and graphic results. The power of this approach is that the bulk of the effort consists in the one-time development of the benchmark model, which remains fixed; variations in the decision aiding scheme may be modeled and analyzed conveniently and efficiently.

The model of the human DM is shown in Petri Net (PN) form [6] in Figure 1. PNs are used here since they permit unambiguous representation of the information structure of an organization. The reader is referred to other works which develop and apply techniques in which PNs yield quantitative results about DM organization performance [7] and organization design [8].

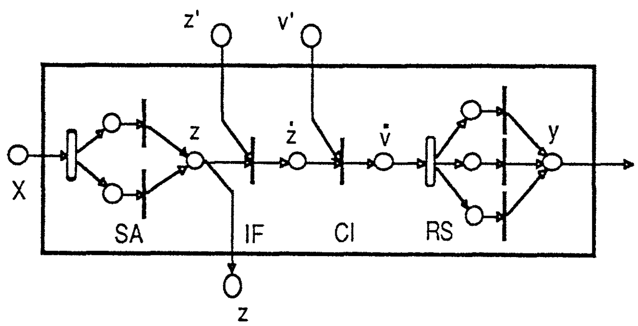

Fig. 1 Petri Net Representation of Decisionmaker D of Organization O

A PN is a bipartite directed multigraph consisting, for the modeling of decisionmaking organizations, of four elements: places, transitions, decision switches, and directed arcs. Places and transitions, respectively, may be thought of as conditions and events. A transition is said to be enabled if every place capable of providing it with input has a token. Tokens are symbolic carriers of information. Firing of an enabled transition removes a token from each input place and generates a token to each of its output places. In the DM model, a decision switch is a transition with more than one output place, the choice of which single place receives a token is specified by a decision rule or strategy.

As seen in Fig. 1, an input signal x, arriving from the environment with average interarrival time t, faces in the generalized DM a four stage process. The first and last of these stages, situation assessment (SA) and response selection (RS), model the actual decisionmaking process, while information fusion (IF) and command interpretation (CI) allow for interaction of the decisionmaker with others in the organization.

* This work was carried out at the MIT Laboratory for Information and Decision Systems with support provided by the Charles Stark Draper Laboratory, Inc. under contract DL-K-260948 and the Joint Directors of Laboratories under contract N00014-K-85-0782 from the Office of Naval Research.

The SA stage consists of a switch and U algorithms (for decisionmaker D, $U = U^d = 1$). The switch may be set by the decision variable, u, according to the internal decision strategy p(u). The selected algorithm, f, operates upon x to produce an assessed situation z. This information may, in turn, be combined with information from other decisionmakers, z', to yield $\bar{z}$.

The fused assessed situation, $\bar{z}$, is processed by one of V RS algorithms ($V^d = 2$). The CI stage of the model allows $\bar{z}$ and the external information v' to influence the choice of this algorithm; v' may be considered to be a command capable, for example, of restricting response options. The RS algorithm, h, is chosen according to a second strategy, $p(\bar{v} \mid \bar{z}, v')$. Note that no restriction is placed on the algorithms themselves, other than that they be well defined on a finite set of internal variables w_i^j, where $i=1,\ldots,n_j$ is the number of variables in algorithm j and $j=1,\ldots,(U+V)$ is the number of SA and RS algorithms in decisionmaker D.

It will be necessary in the application to invoke the concept of a preprocessor. Preprocessors operate between an information source and a decisionmaker. As modeled by Chyen [9], they may describe an external decision aid or an internal subsystem of the decisionmaker. The purpose served by the preprocessor is to influence the internal decision strategy by gaining knowledge about x.

The structure of the DM model provides a convenient framework for the application of information theory [1], [2], [3], [4]. Using entropy, H, as a measure of the uncertainty in a set of random variables, the total information processing activity of a decisionmaker, G, may be expressed as

$$G = \sum_{i=1}^{n} H(w_i) \qquad (1)$$

where the w_i's are all the variables of the decisionmaker defined above. This total workload may be broken down according to the Partition Law of Information of Conant [10] in the following way:

$$G = G_t + G_b + G_c + G_n \qquad (2)$$

where G_t, denoting the throughtput, or information transmission from input to output, is given by

$$G_t = T(x{:}y) = H(y) - H_x(y). \qquad (3)$$

G_c, the coordination, a measure of interconnectedness, may be expressed as

$$G_c = T(u{:}w_1{:}\ldots{:}w_{n-1}{:}\ y); \qquad (4)$$

the noise, G_n, i.e., the entropy remaining in the internal variables and the output y when the input x is fully known, as

$$G_n = H_x(w_1,\ldots,w_n,y); \qquad (5)$$

and finally the blockage, G_b, as

$$G_b = T_y(x{:}w_1{:}\ldots{:}w_{n-1}). \qquad (6)$$

These four quantities do not, nor does this work, take into account information present in the input but rejected by the system. The additional assumptions necessary for the application of information theory in this work are: (1) the model is memoryless; (2) the algorithms are deterministic; (3) the algorithms have no rejection; and (4) the sets of algorithm variables are mutually disjoint, i.e., only one algorithm is active in each stage at any particular time.

As one would expect, and can deduce from equations (2) through (6), the total activity, G, of a DM is a function of the individual's decision strategy or choice of decision algorithms. Because DMs in an organization may interact, an individual's workload is in general a function of the organization decision strategy. In much the same way, the organization performance index, J, is also a function of the organization decision strategy. The precise function J may be specified by the modeler as, for example, a probability of error. As shown in Fig. 2, J is computed as a function of the difference between the actual and desired organization response, where L(x) is the desired x to y mapping, presumed to be known by the modeler. The quantity d(y,Y) is the difference between the actual and desired responses and is presumed to be measurable.

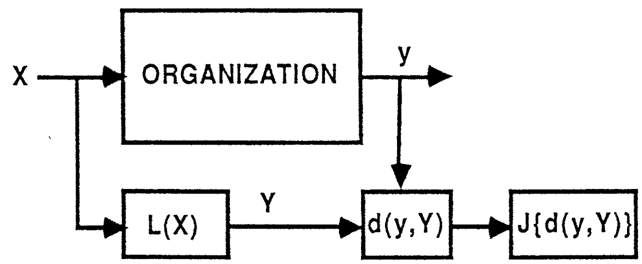

Fig. 2. Computation of the Organization Performance, J

In the mathematical theory of organizations developed in [2], [3], [4], [5] and described here, the kth pure internal decision strategy of DM r, may be defined as

$$D^r_k = \{\delta^r{}_i^{SA} = p(u=i),\ \delta^r{}_j^{RS} = p(\bar{v}=j \mid z = z_m,\ v' = v'_m\} \qquad (7)$$

where: $\delta^r{}_i^{SA}$ and $\delta^r{}_i^{RS}$ are decision strategy parameters corresponding to distributions on SA and RS strategies respectively; $\bar{v}$ is the RS decision variable; $z_m \varepsilon Z$ is a fused assessed situation; and $v'_m \varepsilon V'$ represents a command input. The strategy is known as a **pure** strategy if both probabilities equal one, otherwise it is a **mixed** strategy. For this model of decisionmaking, an upper bound on the number of pure internal decison strategies is given by the expression

$$n_r = U \bullet V^M \qquad (8)$$

where U, V, and M are respectively the number of algorithms in the SA and RS stages of the DM, and the dimension of the set $\bar{z}$. The interaction among decisionmakers means that workload and performance are functions of the strategy of the organization taken as a whole, i.e. its organizational strategy, given by the r-tuple

$$D_{i,j,\ldots,k} = \{D^1_i, D^2_j, \ldots, D^r_k\} \quad (9)$$

where r is the number of DMs in the organization and i,j,...,k are the pure internal strategies defined in (7). By computing J and G as as functions of strategy and plotting J against G parametrically with respect to $D_{i,j,\ldots,k}$ the performance-workload (p-w) locus for the organization can be constructed. This locus serves as the quantitative and qualitative/graphic basis for analysis.

ANALYTICAL APPROACH FOR EVALUATING DECISION AIDS

The approach proposed here for evaluating decision aids is general and straightforward:

Step 1. *Model the organization without a decision aid and compute the individual DM workloads and the organization performance in order to establish a benchmark.* The results of the analysis should be helpful in establishing precisely where and under what conditions the organization is overloaded or suffers from poor performance, presumably what the aid is to correct.

Step 2. *Modify the model of the organization to incorporate the proposed decision aid.* Note that when using this methodology, it is sufficient here to model the information to be presented by the proposed aid and to do so in the same fashion as the task input was modeled. The same methodology used to establish the benchmark measures is employed to compute the descriptive workload and performance characteristics of the aided organization.

Step 3. *Compare the results for the aided organization to the benchmark.* The effects of the decision aid on the workload of the individual user(s) and on the performance of the organization as a whole can be isolated to determine the overall effects of the aid. In addition, the implications of meta-decision can be investigated by examining the results of aiding strategies in either pure or mixed form. The results obtained in this third step are well-suited for normative use, indicating how the aid may be designed or used to bring about improvements.

Step 4. *Steps (2) and (3) may be iterated for comparative evaluation of alternative proposed decision aid designs.*

This approach to evaluating decision aids is intuitive and straightforward. In order to render the specifics of the approach more clear, the following section illustrates, by way of an application, how the tools of mathematical organization theory presented in the introduction are brought to bear in each step of the approach.

APPLICATION: SUBMARINE EMERGENCY DECISIONMAKING

The subject of the application is an actual decisionmaking organization, the Ship Contol Party (SCP) of a U.S. Navy submarine performing emergency control. Since only a single aid will be considered, steps 1 through 3 will be illustrated.The submarine context was chosen for the following reasons: (i) the SCP performing emergency control has a non-trivial, representative, organization structure consisting of five DMs, exhibiting both parallel and hierarchical characteristics, (ii) the SCP members are well-trained for their respective decision tasks, (iii) the SCP, in performing emergency control, has characteristics which make it attractive for information theoretic modeling, (e.g. it has well-defined decision rules and it must make its decision in a matter of seconds, minimizing the likelihood of loops in the organization structure), (iv) the SCP has been studied as a candidate for decision aiding [11].

AN OVERVIEW OF SUBMARINE EMERGENCY DECISIONMAKING

Submarine emergency decisionmaking encompasses all decisions about "actions taken to counteract the effects of any and all system failures which impede the normal operation of the submarine and the accomplishment of its mission." [11] System failures ranging from ones of little consequence to those threatening the loss of the ship arise from a variety of sources, including human error and battle damage, and are magnified by the high speed of modern submarines. With operating depths on the order of only five times the length of the vehicle, a distressed vessel may, within tens of seconds, plunge to crush depth, or ascend to broach the surface, disclosing presence and position to the enemy, exacerbating the casualty, or even colliding with another vessel. Clearly there is a rationale for improving speed and accuracy of the decisionmaking response of submarine crews confronting emergency situations.

The large number of possible emergency situations fall into several classes. Among the most dangerous of these is "loss of control" - specifically of control surfaces. Another highly dangerous class is flooding due to failed pipe or hull penetration seals or to damage inflicted by an external agent. Other classes include fire, loss of power, electrical failure, and indicator failure. The occurence of more that one casualty at a given time is known as a compound casualty. Although the detection and recognition of these emergencies is automated to a small degree aboard modern submarines, the SCP members are trained through drill and supervised experience to perform many of these functions themselves. They must fuse gathered information, and decide upon and effect a response with the ballast tanks, control surfaces (rudder, stern planes, fairwater planes), and propeller. Although a thorough familiarity with, and constant cross-checking of indicator readings improves the chances for early casualty detection, recovery from a casualty depends upon closely coordinated information sharing by the SCP and the processing of information from upwards of fifteen sources, according to complicated decision rules, within a matter of seconds. Fig. 3. depicts the SCP positions before the ship and ballast control panels.

On the ship control panel are indications of ship state (speed, depth, heading, trim, and roll), control surface positions, and control mode buzzer and lights (which indicate when the power to a set of control surfaces has failed). The ballast control panel provides indications about the ship's depth and trim conditions, the status of its ballast tanks and pressurized air banks, as well as other non-weapon system information. This panel is equipped with a telephone for communicating with the other ship compartments about, most importantly in the emergency context, flooding. An additional source of information is a loudspeaker informing the crew of surfaced and submerged sonar contacts and tactical situations with implications for emergency decisionmaking.

Emergency control of U.S. submarines consists of two phases, immediate and supplementary actions [12]. Immediate actions are those which must be performed in seconds, if potentially catastrophic consequences are to be averted. Supplementary actions are follow-up measures for minimizing the effects of a casualty which need not be performed within a

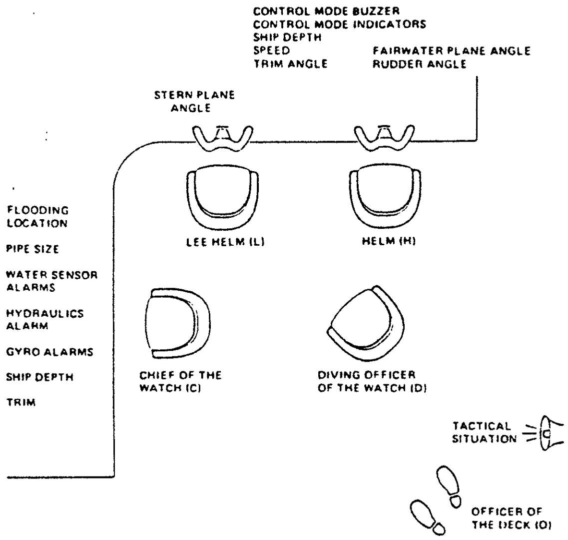

Figure 3. (The Layout of the Ship Control Party and the Ship and Ballast Control Panels)

On the ship control panel are indications of ship state (speed, depth, heading, trim, and roll), control surface positions, and control mode buzzer and lights (which indicate when the power to a set of control surfaces has failed). The ballast control panel provides indications about the ship's depth and trim conditions, the status of its ballast tanks and pressurized air banks, as well as other non-weapon system information. This panel is equipped with a telephone for communicating with the other ship compartments about, most importantly in the emergency context, flooding. An additional source of information is a loudspeaker informing the crew of surfaced and submerged sonar contacts and tactical situations with implications for emergency decisionmaking.

Emergency control of U.S. submarines consists of two phases, immediate and supplementary actions [12]. Immediate actions are those which must be performed in seconds, if potentially catastrophic consequences are to be averted. Supplementary actions are follow-up measures for minimizing the effects of a casualty which need not be performed within a strict time frame.

Step 1 Applied: *Model the organization without the decision aid*

THE TASK MODEL

In order to construct an information theoretic model of submarine emergency decisionmaking it was necessary to make several assumptions about the task: (i) an emergency is a discrete event, one of which may occur at an instant in time (the model does not address evolving situations); (ii) emergencies are considered as repeated but independent events, in order that the average activity over time may be computed (recall that the DM model employed here is memoryless); (iii) only emergency situation assessment and response selection are considered, to the exclusion of detection and response implementation; (iv) the model is limited to considering only the most dangerous casualties (hydraulic failure, control surface jams, flooding) and false alarms; (v) the organization input alphabet is defined to include the maximum manageable subset of instruments and internal telephone reports relevant to the modeled casualties; (vi) the organization topology, the distributions on input alphabet, and the DM algorithms that appear in the model are predicated upon the assumption that only the immediate actions, described above, are modeled. For the complete task model specification see [13]. Task assumptions (iv) and (v) were necessary in order to keep the dimensionality of the problem to a level commensurate with computation on an IBM PC AT. The possibility of false alarm was included in order to capture the cost incurred when an unnecessary (noisy) emergency response compromises the secrecy of the submarine's whereabouts. Task assumption (vi) scopes the problem and in so doing renders the previous assumptions more reasonable. For example, the immediate actions are to be decided upon within 5 to 7 seconds, which corresponds roughly to both the rate at which the DMs are trained to "sample" the inputs, as well as to the assumed repeated event interarrival rate. Additionally, the DMs are trained to counteract one casualty at a time when performing the immediate actions; thus, evolving situations are likely to be treated as a sequence of discrete events, as in assumption (i). Task assumption (vi) will be invoked again when the organization topology is modeled, below.

THE ORGANIZATION MODEL

This development shall begin by describing the organization in overview - which decisionmakers receive what information and how the processed information is shared - and shall culminate in a Petri Net representation of the modeled structure. For a description of the process at a lower level, and for a formulation of the models of the individual DMs at the structural and algorithmic levels, the reader should refer to [13].

As seen in Fig. 4, the ship control party is an organization with hierarchical and parallel characteristics, consisting of five DMs: the Officer of the Deck (symbolized as OOD or O), the Diving Officer of the Watch (DOOW, D), the Chief of the Watch (COW, C), the Lee Helm (L), an the Helm (H).

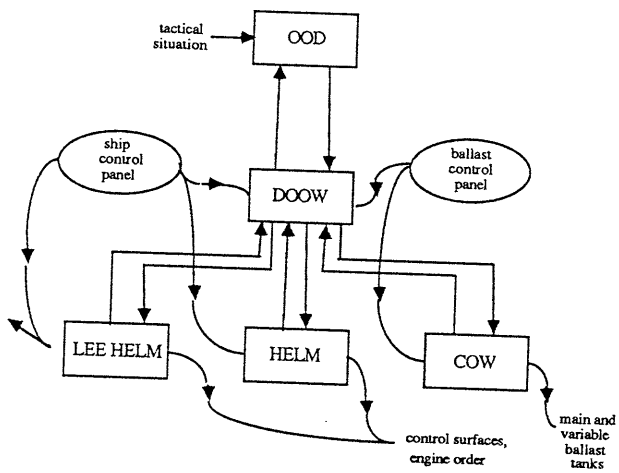

Figure 4. The Ship Control Party

At the top echelon of the hierarchy is the OOD, with the responsibility for integrating the ship control process with the other aspects of the ship's mission. The essence of this job during emergency control is to decide whether certain aspects of the emergency response should be restricted because of the existence of a sensitive tactical situation. Second in command is the DOOW, whose job in the emergency context is to monitor, coordinate, and direct the actions of the sensing and effecting echelon, subject to any restrictions imposed by the OOD. The bottom echelon, operating under commands issued by the DOOW, consists of the COW and the helmsmen. The COW receives all information on flooding casualties and hydraulic failure, which he shares with the DOOW. This DM is also charged with controlling the ship ballast system for aiding in the control of depth. The Lee Helm, L, drives the ship's stern planes, the control surface that modulates the vehicle's trim angle and thus its depth. In carrying out this task, L receives plane angle information, ship state information (speed, depth, trim, etc.). Finally, the Helm, H, controls the ship's rudder and fairwater planes (the small control surfaces locate on either side of the sail, sometimes called conning tower) based on plane angle information, and the same ship state information received by L.

The topology of the modeled ship control party is represented as a Petri Net in Fig. 5. As seen in this figure, the OOD is modeled as a single algorithm, denoted as IF^o, which considers the information fused by the DOOW, $\bar{z}^{do}$, and the tactical situation to produce the command v^o which may restrict the response options available to the DOOW.

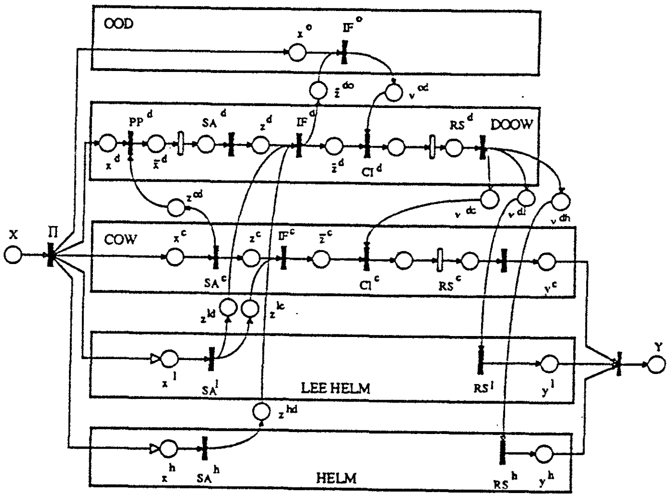

Fig. 5. Petri Net Representation of the Ship Control Party

The DOOW model presents a rich example of an information theoretic DM model. Inputs to the DOOW's preprocessing stage are the partition x^d of the input vector X, as well as shared information from the COW, z^{cd}. While shared situation information normally is fused in the IF stage, the methodology is flexible enough to permit situation information from one DM to be considered in the situation assessment stage of another, as this particular application required.

To apply the methodology rigorously, the following assumptions about the organization were also necessary: (i) the organization structure is fixed in a configuration with a high likelihood of actually occurring, (ii) the organization structure contains no closed loops and processes each discrete task on a single pass. Both of these organization assumptions conform to task assumption (vi), since the extreme time processing constraint imposed by the immediate actions seems to limit the likelihood that the organization structure will evolve over time, and certainly penalizes the tendency to introduce time consuming loops.

With these assumptions made, a model was developed utilizing subjective probabilities and judgements gathered in interviews with a U.S. Navy submariner [14]. The computational implementation is discussed separately below. The results proved robust with respect to variations in the task assumptions [13]; thus the model is appropriate at minimum for purposes of illustration of the approach and for first order substantive guidance. These benchmark results, to be presented with the results of steps 2 and 3, confirmed that one SCP DM in particular, the DOOW, formed an information bottleneck in the organization and was overloaded.

Step 2 Applied: *Model the oganization with the decision aid*

SELECTING THE AID

In this application, to illuminate the modeling and integration of the decision aid, the selection of the aid is discussed. In particular, a number of case-specific concerns must be addressed. First, an SCP decision must be made on the order of seconds, perhaps tens of seconds; because of this extremely constrained time frame, an interactive decision aid would not be feasible. Second, the existing information panel performs a degree of preprocessing, obviating a decision aid with that function. On the other hand, an aid providing an emergency response decision is more in the nature of automation than decision aiding. Presumably, what is desirable is a non-overloaded DOOW in the loop. The only scheme that keeps a man in the loop, satisfies the time constraint, and is not already done by existing instrumentation, would be an aid providing the user with a situation assessment. Therefore, an aid that furnishes the DOOW with an assessed situation shall be evaluated.

IDENTIFICATION AND MODELING OF META-DECISIONS

With the aid functionally specified, it was necessary to make a set of final assumptions in order that the aiding strategies could be defined: (i) the aid will not replace existing information but will be included saliently among the existing instruments, (ii) even with a decision aid available, the DMs will continue to be trained to assess emergency situations unaided, (iii) the aid is absolutely reliable, but the user is not certain about its reliability.

For this application, under assumptions made thus far, the need for the user to make *meta-decisions* emerges. Specifically, the aided DM must, when confronted by an emergency, decide between the following three options: (1) the user DM ignores (blocks) the information provided by the aid and assesses the situation as trained, (2) the user DM assesses the situation as trained and compares the result with aid information, choosing the worse case, (3) the user DM relies solely on the aid information.

The methodology easily captures these meta-decisions with the notion of decision strategy introduced earlier. Whereas the unaided DM's decision strategies represented a choice about which of the available situation assessment decision algorithms to activate, the identical construct allows meta-decisions to be represented as strategies, in this case aiding strategies, where the "algorithms" are functions of the aid information, the unaided situation assessment algorithms, and the output of the unaided algorithms. Referring to Fig. 6, the top aided algorithm consists of the entire (deterministic) preprocessor-situation assessment stages of the unaided decisionmaker, which permit the choice of option (1) above.

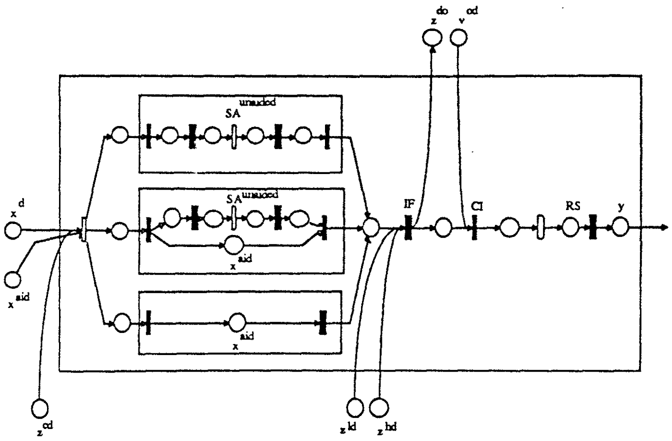

Fig. 6. Petri Net Representation of the Aided Diving Officer of the Watch (DOOW)

The middle algorithm shows the preprocessor-situation assessment in parallel with an identity algorithm for processing the aid's situation assessment, followed by an algorithm for performing worst-case comparison of the two situation assessments as described in option (2) above. The lower algorithm consists simply of the identity algorithm for echoing the decision aid's situation assessment, option (3).

ON COMPUTATIONAL IMPLEMENTATION AND ANALYSIS

Analyses of the models developed in Steps 1 and 2 were performed on an IBM PC AT using a software package written in TURBO PASCAL specifically for this application. Fig.7 schematically depicts the software package modules. First, the task model generates the input vector X, consisting of approximately 1500 possible combinations of the assumed input vector elements, and the probability distribution, p(x), on those states. The next stage consists of simulation models of the unaided and aided organizations in which the modeled decision algorithms, called according to the precedence established in the PN representation (Fig. 5), process all of the letters, X_i, of the alphabet X. While this processing occurs over all X_i, probability distributions are derived for all internal variables w of the system as a function of the decision strategy. Eq.(1) is then applied to compute the average information processing activity of the individual DMs. As illustrated in Fig. 2, the unaided and aided simulation models generate for each X_i a response Y_i. The response Y_i is compared against the known desired response Y_i' according to a cost functional L(Y,Y'), defined here as probability of error, weighted by an index of error gravity, to produce the performance index, J. Finally, the performance-workload locus is computed as a convex function of the organization's decision parameters, defined in Eq. 7.

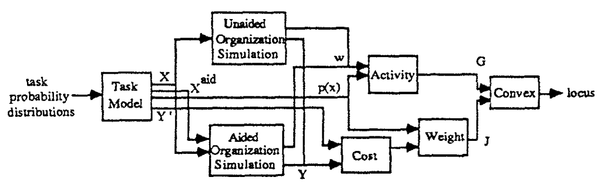

Fig. 7 Computation implementation schematic

Step 3 Applied: *Compare the results for the aided organization against the benchmark*

In Table 1 are presented the results of Step 1, the benchmark (unaided) analysis. Recall that J is the organization performance, an expected error cost, where such costs are defined in the interval [0, 1]. The measure J is dimensionless, and the values shown should be interpreted only as references for comparative analysis. The actual rate of error corresponding to J in Table 1 was approximately 0.10.

Table 1 Results of Step 1: Performance and Workload of the Unaided Organization

	AVERAGE	RANGE	STD. DEV.
J	0.010 - 0.040	0.026	0.001
G^o (bits)	constant	5.140	0.
G^d "	constant	54.05	0.
G^c "	27.21 - 30.55	29.30	0.920
G^l "	constant	11.272	0.
G^h "	constant	7.525	0.

The workload quantities, G^i, measured in bits, also should be used only for comparisons. For example, note how the DOOW workload, G^d, as expected, greatly exceeds that of the other DMs.

The results of Step 2 are given in Table 2. Comparing the results, the most notable effects of the aid are that sizeable percent improvements in performance may result and (2) workload may be either significantly reduced or increased.

Table 2 Results of Step 2: Performance and Workload of the Aided Organization

	RANGE	AVERAGE	STD. DEV
J	0.004-0.040	0.023	0.01
G^o (bits)	5.140-5.141	5.140	0.0004
G^d "	41.143-62.474	54.430	6.020
G^c "	27.258-30.7065	29.354	0.923
G^l "	11.233-11.2715	11.258	0.012
G^h "	7.476-7.525	7.507	0.015

What in fact happens, from one decision to the next, depends upon the outcome of a meta-decision about the use of the aid information. Recall from Step 2 that the meta-decision phenomenon was modeled using the concept of decision strategy. Fig. 8, which depicts a representative slice of the p-w locus for the DOOW such that only the meta-decision is varied, clearly illustrates the impact of meta-decisionmaking.

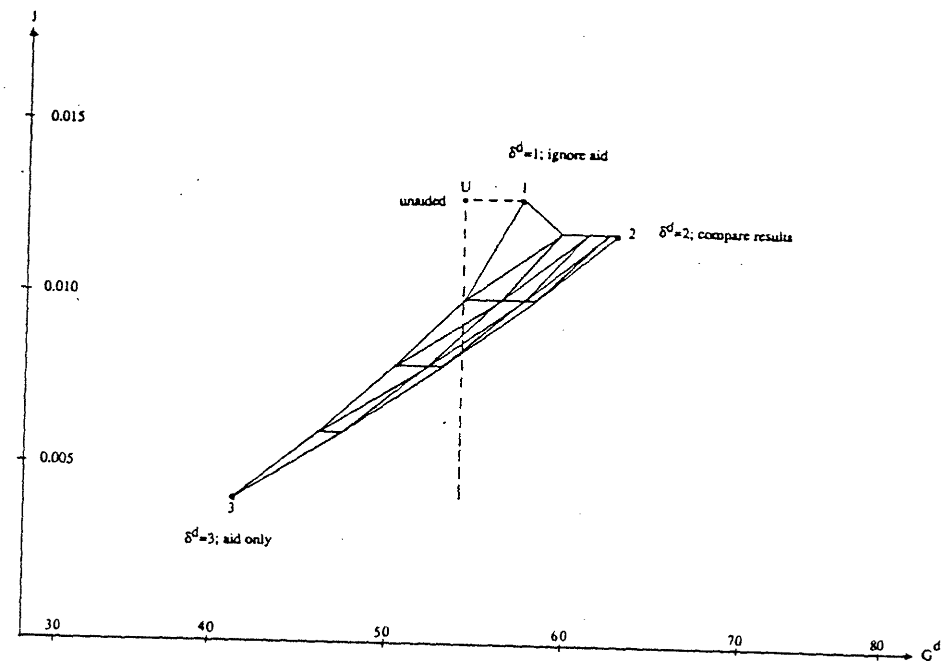

Fig. 8 J versus G^d ($\delta_1{}^c$, $\delta_2{}^c$ held constant)

Listed and labeled are four pure strategy points: U, indicating the benchmark result for that slice; 1, where DOOW, by setting decision parameter δ^d=1, chooses to ignore the aid; 2, where δ^d=2, and DOOW compares results, choosing the worse case; and 3, where δ^d=3, only the aid situation assessment is considered. A pure strategy is one that will be used by a DM with probability unity. The edges and interior of the locus charaterize all mixed strategies, where the probability of a meta-decision selecting any one of the three options is less than unity. Comparing the benchmark point U with point 1, performance is unchanged, as one might expect. A subtle difference, however, is that the workload is slightly higher at point 1. There are two reasons for this increase: (1) an amount of activity is required to suppress, or block, the aid information (Eq. 6); and (2) there is a strictly positive workload associated with meta-decisionmaking, contributed by noise (Eq. 5).

At point 2, where the DOOW would compare a self-generated situation assessment against the aid's and choose the worse case, performance improves slightly (4% on average), however this improvement is accompanied by an approximately 16% rise in workload from 54.05 to 62.47 bits. This increase is attributable to: (1) additional coordination necessary to process two signals, as well as blockage resulting from suppressing the better assessment; and (2) meta-decisionmaking workload, as for point 1. It is necessary to point out that the methodology assumes that DMs are characterized by a bounded rationality constraint, such that exceeding the constraint places the DM in a regime where performance cannot be modeled, and is usually expected to be poor. If the reasonable assumption is made that in an emergency, unaided DOOW processes information as fast as he is able, increasing the emergency decisionmaking workload can be expected to cause the constraint to be approached or violated, risking serious degradation in performance.

Point 3 of Fig. 8 illustrates the expected organization performance and DOOW workload when the aid's assessment, which is assumed to be perfectly reliable and unerring, is always used to the exclusion of self-generated assessment. Here, performance improves by 43% on average (there being no situation assessment errors). At the same time, workload is reduced by 24%, as the elimination of situation assessment processing greatly outweighs the additional workload introduced by meta-decisionmaking.

From these results, we can see that the expected effect of the aid for any individual depends upon where the user's meta-decisionmaking behavior places him in the p-w locus. We can easily compute that, in order to bring about an average improvement in workload and thus performance, the aid-only option must be used at least 50% of the time (which result appears graphically in Fig.8, where the dashed vertical line below point U approximately bisects the locus).

Another interesting result of the comparative analysis in Step 3 is illustrated in Fig. 9. The plot in this figure depicts the p-w locus, from which the "slice" in Fig. 8 was taken. One sees three vertically-oriented crescent-shaped structures which represent the p-w locus when meta-decision strategy is held constant at each pure strategy. Connecting these three sub-loci are surfaces or "slices" created when all organization strategies except the meta-decisionmaking strategies are held constant. Note the resemblance of the slices to those shown in Fig. 8. The significant result is that the variation in performance (the vertical axis) in each slice is small compared to the range, in this dimension, of the locus as a whole. In other words, while for any given organization strategy the aid may have benefits, as shown in the discussion of Fig. 8, the magnitude of these benefits may be wiped out by poor selections of the organization strategy. Even if a favorable aiding strategy is employed, individual decision strategies chosen elsewhere in the organization may determine the quality of the organization response.

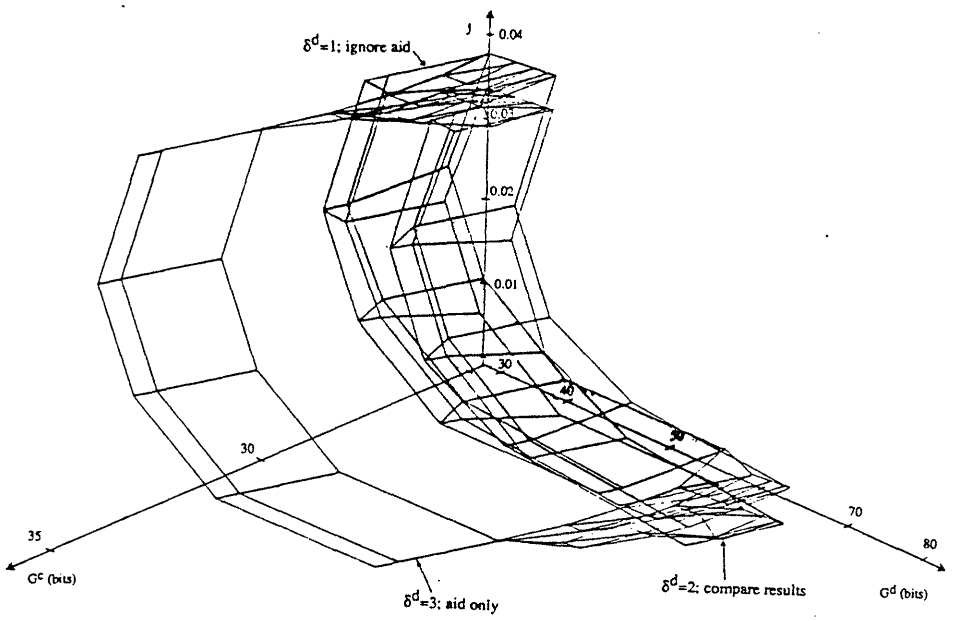

Fig. 9 The pure decision-aiding locus ($\delta_1{}^c$, $\delta_2{}^c$ held constant at pure strategies)

CONCLUSIONS

An analytical approach for the evaluation of decision aids has been presented, that has special relevance to decisionmaking organizations. The steps of the approach were applied in the context of submarine emergency decisionmaking. Results of the application provide evidence that organization context and the phenomenon of meta-decisionmaking play an important role in determining the ultimate effectiveness of a decision aid, and that these matters should be taken into consideration, not only before designing a decision aid, but defore deciding to design one.

REFERENCES

[1] Boettcher, K.L., "An Information Theoretic Model of the Decisionmaker," S.M. Thesis, Report LIDS-TH-1096, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, 1981.

[2] Boettcher, K.L. and A.H. Levis, "Modeling the Interacting Decisionmaker with Bounded Rationality," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-12, No.3, 1982.

[3] Boettcher, K.L. and A.H. Levis, "Modeling and Analysis of Teams ofInteracting Decisionmakers with Bounded Rationality," Automatica, Vol.19, No.6, pp. 703-709, 1983.

[4] Levis, A.H. and K.L. Boettcher, "Decisionmaking Organizations with Acyclical Information Structures," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-13, No. 3, pp. 384-391, 1983.

[5] Levis, A.H., "Information processing and decisionmaking organizations: A mathematical description," Large Scale Systems 7, pp. 151-163, 1984.

[6] Tabak, D. and A.H. Levis, "Petri Net Representation of Decision Models," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-15, No.6, pp. 812-818, 1985.

[7] Hillion, H., "Performance Evaluation of Decisionmaking Organizations Using Timed Petri Nets," SM Thesis, LIDS-TH-1590, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, 1986.

[8] Remy, P., "On the Generation of Organizational Architectures Using Petri Nets," SM Thesis, LIDS-TH-1630, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, 1986.

[9] Chyen, G.H.-L., "Information Theoretic Models of Preprocessors and Decision Aids," SM Thesis, LIDS-TH-1377, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, 1984.

[10] Conant, R.C., "Laws of Information Which Govern Systems," IEEE Transactions of Systems, Man and Cybernetics, Vol. SMC-6, No.4, 1976.

[11] Motyka, P., "An Integrated, Responsive Fault and Damage Tolerant Approach to Submarine Emergency Control," Report CSDL-P-1800, The Charles Stark Draper Laboratory, Inc., Cambridge, MA, 1983.

[12] Anonymous, SSN 688 Class Ship Systems Manual, Vol. 7, Principles of Ship Control, NAVSEA No. S9SSN-W4-SSM-820/ (C) 688 Class V7.

[13] Weingaertner, S.T. "A Model of Submarine Emergency Decisionmaking and Decision Aiding," SM Thesis, LIDS-TH-1612, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, 1986.

[14] Labak, S.J., Private Communication, 1985-6.