Organizational Designs Suited to High Performance Under Stress

Kathleen M. Carley and Zhiang Lin

Abstract—**This paper discusses the role of organizational design in affecting organizational performance. Using a computational framework, CORP, various aspects of organizational design are examined; e.g., training, communication and command structure, and resource access. Organizational performance is evaluated under both optimal operating conditions and stressed conditions (such as information errors and turnover), given both a simple and a complex task. These analyses suggest that when the organization is facing a choice task where all options are equally likely then the best performing organizational structures are simple (i.e., operational training, teams, and lack of overlap in access to information). In contrast, when the task is biased, i.e., one outcome is more likely than another, then the best performing organizational structures are more complex (i.e., training that relies on experience, hierarchies, and overlap in information access).**

I. Introduction

MOST studies of organizational design consider only optimal conditions. However, organizations are frequently stressed. This stress can be due to internal ambiguities such as missing information that result in uncertainty and so can alter what decisions are made [6]. Or, this stress can be due to external critical events such as a hostile aircraft entering the airspace, where the impact of the decision is critical (making the wrong decision may be disastrous) [14]. Regardless of its source, stress degrades performance. However, there has been little effort directed at determining how to design organizations for high performance under stress conditions [3]. Rather much of the work on organizational design has focused on the general impact of design [9], [18], [13].

This paper discusses the application of the CORP framework to designing organizations suited to stress. Previous research has shown that the CORP framework can be used to model and contrast the performance of organizations with a variety of designs [2], [4]. These previous analyses suggest that an advantage of the hierarchical form is that it is minimally affected by internal stress [3], [4]. Whether, in fact, the hierarchy or another design is the better performer under external stress has not been addressed.

Manuscript received August 1, 1993; revised February 24, 1994. This work was supported in part by Grant No. N00014-90-J-1664 from the Office of Naval Research (ONR),

K. M. Carley is with the Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 USA.

Z. Lin is with the H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA 15213 USA.

IEEE Log Number 9405719.

II. CORP

The CORP framework is a simulation testbed designed to allow the researcher to compare and contrast the performance of organizations with different designs subject to different types of operating conditions. CORP uses a choice task similar to that used by many researchers interested in organizational design [10], [11], [17]. CORP models are artificial organizations composed of complex adaptive agents with task specific abilities who are capable of making decisions either on the basis of personal experience or standard operating procedures. It is not the purpose of this paper to describe the CORP framework in detail. Further, details on the instantiation of the CORP framework used in this paper can be found in [5]. In this paper, we briefly highlight those aspects of CORP used in our analysis. The purpose of this paper is to examine the interplay between organizational design and performance under stress. This examination is carried out using CORP.

In the CORP framework the organization is subjected to a series of distributed decision making problems such that each problem is similar to (but not identical to) previous problems[1]. The problems are sufficiently complex that no one agent has access to all the information necessary to make the decision. Decisions are made by integrating decisions made by distributed agents on different aspects of the task rather than by consensus [1]. In the CORP framework each problem requires the organization to choose between a set of choices. In order to make this choice, information must be gathered and evaluated by agents at the lowest level in the organization (the analysts). These agents then pass on their decisions to their superiors who integrate this information and then make their own decision which, depending on the organizational structure, is either passed on to their superiors or is the organization's decision. After the organization makes a decision (selects one of the choices) a new problem is given to the organization, and the process starts over with the analysts gathering information. The organization faces only a single problem at a time. Problems are chosen randomly without replacement from the set of problems in the problem space.

Within this generic framework the researcher can examine numerous variations. The task can be varied across a number of choice tasks. The organizational structure defining who communicates to and commands whom can be altered. Information gathering can be affecting by altering the resource access structure which defines who has access to what information.

[1] Tasks such as these have been referred to as quasi-repetitive integrated decision making tasks [4].

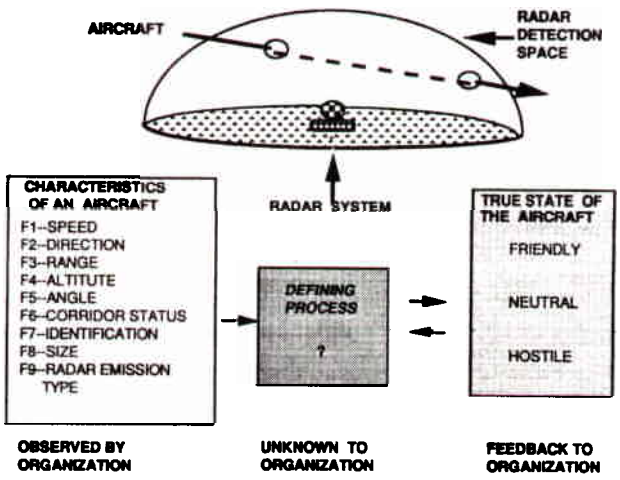

Fig. 1. Stylized radar task.

What decisions agents make can be altered by changing the way in which they are trained. As these variations are made the researcher can examine the impact of the change on organizational performance.

A. Task

The CORP framework can be used with different fixed choice tasks. CORP models have been used to explore both binary choice [2]–[4] and trinary choice [5] tasks where the agents are experientially trained. These choice tasks can be unbiased or biased. In an unbiased task all possible choices are equally likely across the set of problems faced by the organization [4]. In a biased task one of the choices is more likely than the others across the set of problems faced by the organization [2].

In this paper, we focus on the trinary choice task. We consider the behavior of a set of organizations given both an unbiased task and a biased task. This task can be thought of as a stylized radar task (see Fig. 1). That is, each problem is a single object in the airspace. The organization must decide for each object whether it is friendly (=1), neutral (=2), or hostile (=3). Each object is characterized by its score on nine factors, such as speed and weapons type. For each factor the score for that factor is either 1, 2, or 3. The meaning of these scores depends on the factor. For example, if the factor is speed then a 1 would indicate a low speed, a 2 a moderate speed, and a 3 a high speed. Each object is uniquely characterized by a 9 digit number such that each digit is either 1, 2, or 3.

Each object also has a true state. The object is either friendly, neutral, or hostile. The relationship between the object's characteristics (the 9 digit number) and its true state depends on the type of choice task being examined by the researcher. In this paper we examine an unbiased and a biased decomposable task. In both tasks, the true state of the object depends on the unweighted sum of the scores for the 9 characteristics. For example, if the object is 123111333 then the sum is 18. Given this sum, a procedure is used for defining which object is friendly, neutral, or hostile. For the unbiased task the procedure used is: (1) if the sum is ≤ 16 the object is friendly; (2) if the sum is ≥ 20 then the object is hostile; and (3) if the sum is >16 and <20 then the object is neutral. In contrast, for the unbiased task the procedures used is: (1) if the sum is ≤ 13 the object is friendly; (2) if the sum is ≥ 18 then the object is hostile; and (3) if the sum is >13 and <18 then the object is neutral. In the unbiased task approximately 1/3 of the objects are friendly, neutral, or hostile. In the biased task, more of the objects are hostile.

The relationship between the object's characteristics (the 9 digit number) and its true state is not known apriori by the organization or any agent within the organization. Rather, over time, the agents and the organization must discover this relationship.

B. Organizational Designs

A variety of factors contribute to organizational design [9], [18]. Among those factors typically identified as important to organizational performance are: the formal structure (who commands or communicates to whom—C2), the resource access structure (who has access to which resources or incoming information), and the procedures for training personnel to make decisions.

For each factor we consider several variations that commonly appear in the literature. We examine three organizational structures—team, hierarchy, and matrix—shown in Fig. 2. These structures vary in communication complexity with the team being the most simple and the matrix the most complex. These structures vary in the level of command—one in the team and two in the hierarchy or matrix — and in the number of managers an analyst must report to—one in the team and hierarchy and two in the matrix. In all structures, the organizational decision is made by the commander. In all structures, all personnel make decisions on the basis of the information available to them and pass on these decisions to their immediate superior(s). In the hierarchy and matrix the information available to the mid-level personnel is the set of decisions made by their subordinates. In all structures, the only agents having access to "raw" information on the problem (rather than the decisions of other agents) are the agents at the lowest level, the analysts. In all structures, the organization's decision is the design made by the top-level manager.

We consider three resource access structures segregated, blocked, and distributed—shown in Fig. 3. The resource access structure defines which analyst has access to which incoming information and the resources to collect that information. The structures shown in Fig. 3 vary in complexity with the segregated being the most simple and the distributed the most complex. In the segregated structure, each analyst has access to different information and thus a completely different view of the problem. In contrast, in both the blocked and the distributed structures each analyst shares some information with other analysts. In the blocked structure, three analysts share exactly the same information and so have identical views of the problem. Whereas, in the distributed structure while no two analysts see exactly the same information each analyst shares some information with two other analysts. As a final point, these resource structures are designed such that if the organization is a hierarchy using blocked access, all three analysts in the block report to the same person.

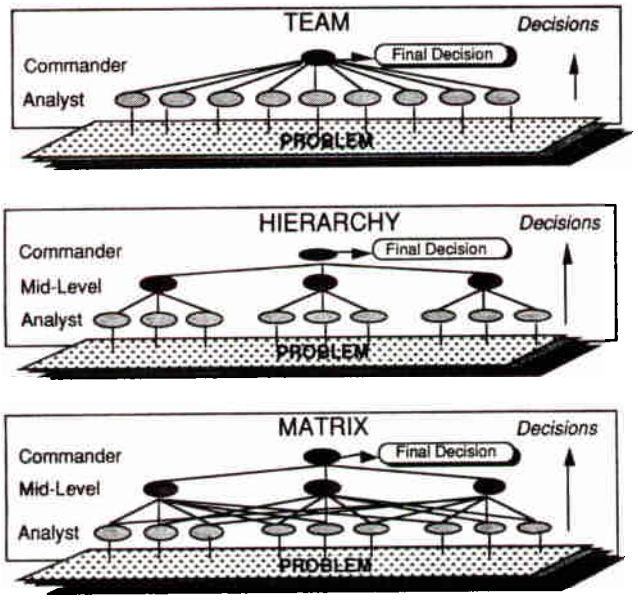

Fig. 2. Stylized organizational structures.

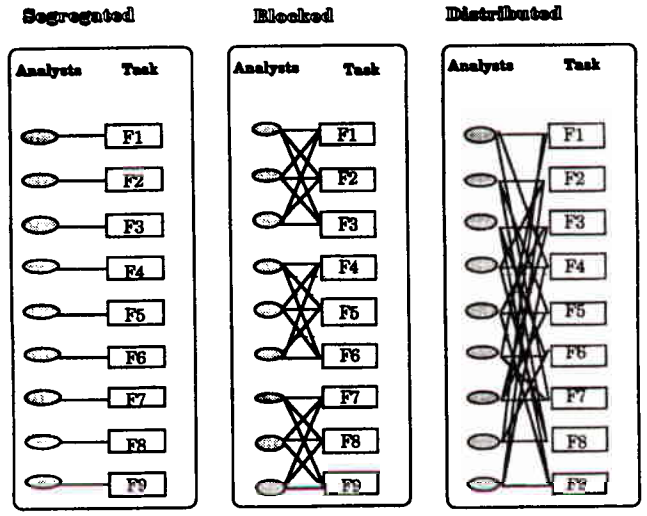

Fig. 3. Stylized resource access structures.

Finally, we consider two training procedures—experiential and operational. In the experiential procedure, each agent in the organization is trained by being given a sequence of all possible problems and being told, after making a decision on a problem, what the correct choice would have been. This training allows the agents to develop a set of procedures specialized relative to their personal experience. The agents can then use these specialized procedures to make decisions. In forming these procedures the agents are told for each problem what was the correct choice for the entire problem. Whether and how this feedback is used depends on the agent's training.

Experientially trained agents use this feedback. Their memories contain an exhaustive list of all incoming patterns of information that they will see and for each pattern the percentage of problems they have had experience with that corresponded to each of the three choices. As the agent gains experience these percentages change. For example, if the agent examines three pieces of information for a trinary task then the agent knows information for 27 patterns (including $111, 112, \ldots, 333$). Experientially trained agents have accurate memories vis these percentages. Agent memory contains information like "for pattern 132 25% of the time the object was friendly, 50% of the time the object was neutral, and 25% of the time the object was hostile." When experiential agents need to make a decision on a new problem they simply locate these stored percentages for the observed pattern, determine which choice (friendly, neutral, or hostile) was most frequently associated with this pattern in the past, and take this choice as their decision. Agents trained experientially act as though they are following a series of rules that adapt over time.

Operationally trained agents do not use this feedback. Operational agents are trained to follow a set of standard operating procedures (SOPs). This operational training allows the agent to employ a set of general purpose procedures that do not adapt over time. Agents trained operationally are assumed to have perfect memory for the SOPs and to execute them perfectly every time. These SOPs are based on the sum of the values associated with those pieces of information available to the agent. Recall that each agent, regardless of its position in the organization sees a set of information such that each piece can take on the value 1, 2, or 3. When the agents handle three pieces of information, as do those at the lowest level in the organization, the SOPs are: (1) if the sum is ≤ 5 then report friendly, (2) if the sum is ≥ 7 then report hostile, and (3) if the sum is >5 and <7 report neutral.

By varying these aspects of design we are able to consider 18 different organizational designs. We examine the performance of these organizations under both optimal conditions and stressed conditions.

C. Stress

We consider two types of stress—internal and external. Internal stress occurs when things go wrong in the organization. As noted by Cohen et al. [7], organizations are messy and things continually go wrong such as people being ill or phonelines breaking. We consider a variety of internal stressors each of which produces uncertainty: missing information, incorrect information, agent unavailability, communication breakdowns, and agent turnover. These stresses occur at the lowest level in the organization. Thus it is the "raw" information on the object's characteristics that is either missing or incorrect. This affects the analyst's decision and so that of his or her superior. Analysts are also the only agents who become unavailable. Further, if there is a communication breakdown the analyst cannot contact his or her superior(s). Note, agent unavailability and communication breakdowns are effectively the same for the team and hierarchy, but have different effects in the matrix. Finally analysts are the only agents that turn over. When agent turnover occurs experiential analysts are replaced with experiential analysts who have no experience; whereas, operational analysts are replaced with operational analysts who have no experience but have been told the SOPs. Only one type of internal stress occurs at a time, and there are three incidents at a time (e.g., 3 pieces of information may be missing).

External stress occurs when the organization is faced with a problem for which a wrong decision is highly costly. For the stylized task we described we define external stress as occurring when the object is hostile. For the unbiased task 33% of the problems are hostile and represent stress events. In the biased task 58% of the problems are hostile.

D. Measuring Performance

Organizational performance is measured in several ways[2]. First, we consider the organization's accuracy. The organization's accuracy is measured as the percentage of problems of a certain type for which the organization's decision matches the true state of the object. For example, we can look at overall accuracy or accuracy under stress. Overall accuracy is measured as the percentage of the 19683 problems for which the organization determined the correct choice when it was operating under optimal conditions (no internal stress). Accuracy under internal stress is measured as the percentage of the 19683 problems for which the organization determined the correct choice when it was operating under a severe level of internal stress (three errors). Accuracy under external stress is the percentage of the hostile cases for which the organization made the correct choice given that it was operating under optimal internal conditions. For the unbiased task this percentage is based on 6561 problems and for the biased task on 11411 problems.

Second we consider the severity of the organization's fallibility. This is measured as the percentage of all mistakes made by the organization that are severe. A severe mistake occurs when the organization claims as friendly an object that is hostile or vice versa. In contrast, a less severe mistake would occur if the organization claimed as neutral an object that was friendly or hostile, or claimed as friendly or hostile an object that was neutral. As with accuracy, this severity of fallibility measure can be measured across all problems or some subset.

E. Relative Performance

We now consider what type of organizational design exhibits high performance under stress. We do this, by examining the relative performance of those designs we have identified. First, however, it is worth noting that regardless of the type of stress, operationally trained organizations tend to be most accurate in unbiased task environments and experientially trained organizations tend to be most accurate in a biased task environments. Further, regardless of the type of stress, operationally trained organizations tend to make fewer severe errors in an unbiased task environment and experientially trained organizations tend to make fewer severe errors in biased task environments. We refer to this as the training-task effect. Essentially, this effect is due to experiential training allowing agents to adapt to the peculiarities of a biased environment. However, this same adaptation can be "over-adaptation" or making mountains-out-of-mole-hills in an unbiased environment.

[2] An elegant method for presenting performance for binary choice tasks is to use ROC and TROC curves. However, we are using a trinary choice task and therefore ROC and TROC curves are not appropriate.

TABLE I
ACCURACY UNDER INTERNAL STRESS

Training Type	Experiential		Operational	
Task Environment	Unbiased	Biased	Unbiased	Biased
Type of Stress				
Missing Information	47.57	59.64	63.57	45.68
Incorrect Information	42.66	59.71	54.61	42.66
Agent Unavailability	47.40	62.21	68.52	46.71
Communication Breakdown	47.41	61.90	69.68	46.74
Agent Turnover	43.87	48.10	76.22	46.23

Note: The Number of cases in each cell is 27.

Now, let us consider the role of internal stress (see Table I). In Table I the accuracy of organizations each type of internal stress is shown. Accuracy for each organization is measured across all 19683 problems. Other than the training-task effect we see that agent unavailability and communication breakdowns are generally less detrimental to organizational performance than are missing or incorrect information. This suggests that managers, by virtue of their position and not necessarily their expertise, are better at adjusting for information errors than are analysts. Further, we see that agent turnover is only detrimental to the experiential organization due to the extra length of time required to train agents experientially than operationally.

Internal stress also affects the severity of the errors made by the organization (Table II). For example, experiential organizations may be almost equally accurate when faced with either missing or incorrect information. However, they are more likely to make severe errors when the information is inaccurate (rather than missing). Experiential organizations are much more likely to make severe errors when beset by turnover than by other forms of internal stress. If the environment is biased, experiential organizations make many fewer severe errors (unless faced with turnover) than do operational organizations (under either biased or unbiased environments). However, if the environment is unbiased, experiential organizations tend to make more severe errors than do the operational organizations (under either biased or unbiased environments).

Under internal stress, the type of stress determines what organizational design fairs best (see Table III). In Table III the organizations who are the most and the least accurate given a severe level of stress (three occurrences) for all 19683 possible problems. For unbiased tasks, teams or hierarchies with segregated access structures are often the most accurate; whereas, for biased tasks, hierarchies or matrices with distributed access structures are typically best. However, this generalization is not completely accurate and we see striking dissimilarities in performance depending on the type of training. For example, the hierarchy with blocked access is the best when there are communication breakdowns and the task is unbiased if the organization is trained operationally,

TABLE II
FALLIBILITY UNDER INTERNAL STRESS

Training Type	Experiential		Operational	
Task Environment	Unbiased	Biased	Unbiased	Biased
Type of Stress				
Missing Information	26.44	4.29	13.64	17.47
Incorrect Information	30.48	5.58	22.41	24.14
Agent Unavailability	22.22	2.78	7.97	13.79
Communication Breakdown	21.77	2.79	6.66	12.24
Agent Turnover	26.85	19.19	4.37	12.26

Note: The Number of cases in each cell is 27.

TABLE III
ACCURACY OF ORGANIZATIONS UNDER SEVERE INTERNAL STRESS, TOP AND BOTTOM PERFORMERS

Training Type		Experiential		Operational	
Task Environment		Unbiased	Biased	Unbiased	Biased
Type of Stress					
Missing Information	T	Hierarchy Distributed 53.51	Matrix Distributed 64.29	Team Distributed 62.45	Hierarchy Distributed 46.20
	B	Team Segregated 39.18	Matrix Segregated 56.26	Matrix Segregated 53.42	Matrix Segregated 41.03
Incorrect Information	T	Hierarchy Segregated 43.42	Matrix Distributed 61.33	Team Segregated 48.83	Team Segregated 41.97
	B	Hierarchy Blocked 35.15	Hierarchy Segregated 56.39	Hierarchy Blocked 47.52	Matrix Blocked 37.64
Agent Availability	T	Hierarchy Segregated 53.91	Hierarchy Distributed 72.64	Hierarchy Blocked 75.85	Team Distributed 47.54
	B	Hierarchy Distributed 35.73	Hierarchy Segregated 57.23	Matrix Segregated 59.59	Matrix Segregated 44.63
Communication Breakdown	T	Hierarchy Segregated 54.84	Hierarchy Distributed 72.60	Hierarchy Blocked 75.82	Matrix Distributed 49.94
	B	Hierarchy Blocked 35.76	Hierarchy Segregated 57.21	Hierarchy Segregated 61.96	Hierarchy Blocked 46.82
Agent Turnover	T	Matrix Distributed 46.47	Hierarchy Distributed 51.60	Team Segregated 100.00	Hierarchy Distributed 47.71
	B	Team Blocked 36.01	Matrix Blocked 30.03	Matrix Blocked 65.26	Team Segregated 43.69

but, if it is trained experientially it is the worst structure. Operational organizations can often do better with slightly simpler structures; e.g., team instead of hierarchy, hierarchy instead of matrix, segregated instead of distributed. Whereas, experiential organizations can take advantage of the more complex organizational forms.

Organizations who are the most accurate are not necessarily the least likely to make severe errors (see Table IV). When the task is unbiased, those operational organizations that are the most accurate also make the fewest severe errors. In contrast, the experiential organizations who are the most accurate are rarely the least fallible. For experiential organizations there is an interesting interaction between type of stress, accuracy, and fallibility. Experiential organizations who face stress due to agents (unavailability or communication breakdowns) are most likely to be accurate if they are organized as a hierarchy with a segregated structure and least fallible vis severe errors if they are organized as a team with distributed structure. In this case, organizational learning as embodied in a more detailed chain of command increases the ability of the organization to be right, but the added uncertainty absorption also increases the chances that the errors that are made are severe. In contrast, if the organization is subject to high information errors then the less complex hierarchy is more accurate but the more costly (in terms of communication channels) matrix is less fallible. Increased communication channels for the experiential organization decrease the likelihood that the organization will learn general procedures that will enable improve overall accuracy and they increase the likelihood that the organization will learn specific procedures that will decrease fallibility. Thus, increased communication channels can lead to over-specialized learning at the organizational level.

When the task is biased, those experiential organizations that are the most accurate also make the fewest severe errors. In contrast, the operational organizations who are the most accurate are rarely the least fallible. For operational organiza-

TABLE IV
FALLIBILITY OF ORGANIZATIONS UNDER SEVERE INTERNAL STRESS, TOP AND BOTTOM PERFORMERS

Training Type		Experiential		Operational	
Task Environment		Unbiased	Biased	Unbiased	Biased
Type of Stress					
Missing Information	T	Matrix Blocked 18.99	Matrix Distributed 2.38	Hierarchy Blocked 10.85	Team Blocked 14.96
	B	Hierarchy Segregated 41.87	Hierarchy Segregated 7.39	Matrix Segregated 26.00	Matrix Segregated 26.93
Incorrect Information	T	Matrix Blocked 25.44	Matrix Distributed 4.99	Team Segregated 19.74	Team Segregated 20.83
	B	Hierarchy Segregated 44.92	Team Segregated 8.91	Matrix Blocked 37.75	Matrix Blocked 34.73
Agent Availability	T	Team Distributed 11.87	Hierarchy Distributed 0.26	Hierarchy Blocked 3.73	Hierarchy Blocked 5.33
	B	Hierarchy Segregated 39.90	Hierarchy Segregated 6.33	Matrix Segregated 17.50	Matrix Segregated 23.14
Communication Breakdown	T	Team Distributed 11.78	Hierarchy Distributed 0.22	Hierarchy Blocked 3.80	Hierarchy Blocked 5.29
	B	Hierarchy Blocked 39.89	Hierarchy Segregated 6.38	Team Segregated 10.90	Matrix Segregated 17.66
Agent Turnover	T	Hierarchy Blocked 22.23	Hierarchy Segregated 12.02	Team Segregated 0.00	Team Segregated 0.00
	B	Hierarchy Segregated 42.49	Matrix Blocked 34.50	Matrix Blocked 12.03	Matrix Blocked 19.68

TABLE V
ACCURACY OF ORGANIZATIONS UNDER EXTERNAL STRESS, TOP AND BOTTOM PERFORMERS

External Condition		External Stress		Overall	
Training Type		Experiential	Operational	Experiential	Operational
Task Environment					
Unbiased	T	Hierarchy Distributed 99.86	Team Segregated 100.00	Matrix Distributed 52.97	Team Segregated 100.00
	B	Team Segregated 51.84	Hierarchy Blocked 88.53	Hierarchy Blocked 34.91	Matrix Blocked 65.26
Biased	T	All Blocked 100.00	Matrix Blocked 78.64	Matrix Distributed 68.42	Hierarchy Distributed 47.71
	B	Team Segregated 96.53	Team Segregated 57.30	Hierarchy Segregated 56.57	Team Segregated 43.69

Note: The number of cases in each cell is 5.

TABLE VI
FALLIBILITY OF ORGANIZATIONS UNDER EXTERNAL STRESS, TOP AND BOTTOM PERFORMERS

External Condition		External Stress		Overall	
Training Type		Experiential	Operational	Experiential	Operational
Task Environment					
Unbiased	T	Hierarchy Blocked 0.00	Team Segregated 0.00	Matrix Blocked 10.2	Team Segregated 0.00
	B	Team Segregated 45.23	Matrix Distributed 39.68	Hierarchy Segregated 37.53	Matrix Blocked 12.03
Biased	T	Except Team All 0.00	Team Segregated 0.00	Matrix Distributed 0.28	Team Segregated 0.00
	B	Team Distributed 48.16	Matrix Blocked 86.14	Hierarchy Segregated 7.04	Matrix Blocked 19.68

Note: The number of cases in each cell is 5.

tions there is an interesting interaction between type of stress, accuracy, and fallibility. In the biased task, the organizational structure that is most accurate if the organization is trained operationally varies by type of internal stress. However, the operational organization will make the fewest severe errors if it is organized as a team when there are information errors and as a hierarchy when there are agent errors.

Now let us consider external stress (see Table V). Accuracy is now based on only that portion of the problems that are considered hostile. All organizations are operating under optimal internal conditions (i.e., no internal stress). We see that no one design dominates. In fact, what is best under one condition may be worst under another. For example, the team with the segregated access structure is the most accurate when the organization is operationally trained and the task is unbiased, and otherwise it is the least accurate structure. It is also the case that regardless of their accuracy, for this task many organizations are not severely fallible when faced only by an external stress (see Table VI).

There are interesting interactions between stress, accuracy, and fallibility. When the task is unbiased and the organization faces external stress the operational organization makes every decision correctly if it is organized as a segregated team. This simply shows that it is possible to locate an optimal structure for the task. In addition, we see that the experiential organization that is distributed may be the most accurate but it will make severe mistakes. Whereas, the experiential organization using a blocked access structure makes more mistakes but none of them are severe. For the experiential organization when faced with an unbiased task performance is affected more by who has access to what resources than by who communicates to and commands whom. When the task is biased, it is the operational organizations who have one structure that is the most accurate (blocked matrix) and another that is the least severely fallible (segregated team).

These results seem to suggest that the organization should simply choose that structure appropriate to the task and the type of training they wish to employ. However, were this to be done, the organization would not necessarily benefit. First the structure that is the most accurate under external stress is not necessarily the most accurate overall. However, the structure that is the least fallible under external stress may be the least fallible overall. For example, for operational organizations teams with segregated access structures simply do not make severe errors (under optimal internal operating conditions).

As can be seen in Tables V and VI, for operationally trained organizations the top and bottom structures are the same whether one is considering performance overall or under external stress. This is not the case for the experientially trained organization. Essentially, operational training is so rigid and has such a compelling effect on performance that the interaction between structure and operating condition is regularized. This is consistent with theoretical arguments about the use of SOPs to buffer the organization [18].

TABLE VII
PEARSON CORRELATION COEFFICIENTS FOR PERFORMANCE

Factor	Accuracy Trinary Task	Fallibility Trinary Task	Accuracy Binary Task
Task Environment	0.246 ***	0.074*	0.153 ***
Training	0.479***	-0.309***	0.586***
Organizational Structure	0.180***	0.186***	0.157***
Resource Access	0.142***	-0.042***	0.036
Internal Stress	0.008	-0.067	0.183***
Internal Stress excluding turnover	0.151***	-0.110**	0.266***
Level of Internal Stress	-0.196***	0.134***	-0.196***

Note: All organizations are facing external stress.
* = .05
** = .01
*** =.005

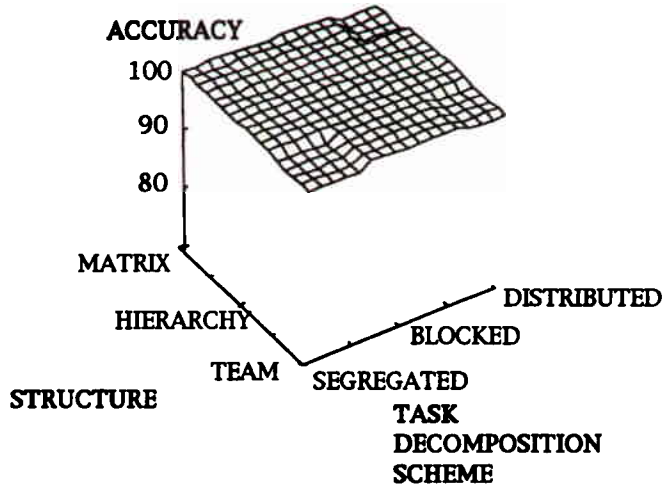

Fig. 4. Average accuracy of experientially trained organizations under external and internal stress given a biased task (percentage of correct decisions/all decisions).

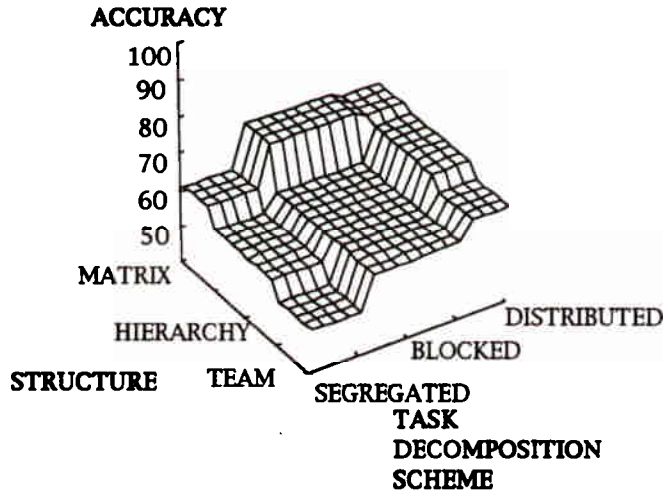

Fig. 5. Average accuracy of operationally trained organizations under external and internal stress given a biased task (percentage of correct decisions/all decisions).

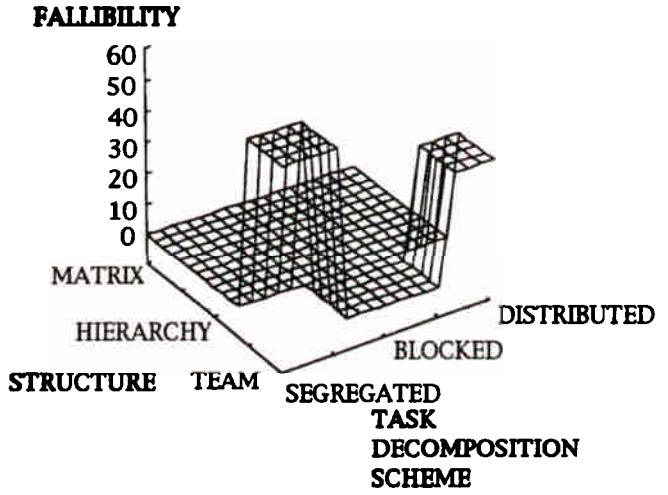

Fig. 6. Average fallibility of experientially trained organizations under external and internal stress given a biased task (percentage of severe errors/all errors).

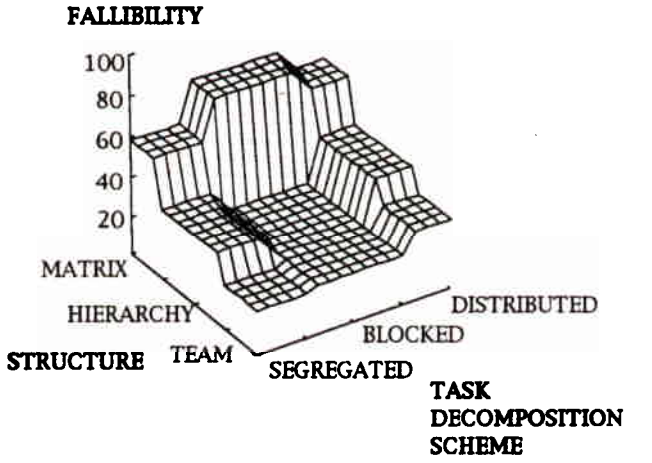

Fig. 7. Average fallibility of operationally trained organizations under external and internal stress given a biased task (percentage of severe errors/all errors).

Finally let us consider the situation typical of organizations operating in the corporate environment. In the corporate world organizations typically operate in a biased environment and simultaneously must contend with both internal and external stress. Organizations using experiential training fair quite well given this scenario and have the luxury of choosing among many organizational structures (see Fig. 4). In contrast, for the organizations using operational training there are sharp differences in accuracy due to organizational design (see Fig. 5). Under these conditions organizations tend to do best, regardless of their training procedure if they are designed as a matrix with a distributed resources access structure and worst if they are designed as a team with a segregated resource access structure.

In terms of fallibility, we see a similar pattern. For experientially trained organizations most organizations are not fallible (see Fig. 6). However, teams are highly fallible. In contrast, under operational training teams are relatively infallible (see Fig. 7). More complex structures, such as matrix structures, are more fallible.

In summary, we find that there is no single organizational design that is best under all stress conditions. The general relationship between the various factors we have considered and performance is summarized in Table VII as a series of correlations for the main effects. We find that organizations tend to be more accurate if a) the task environment is unbiased, b) the members of the organization are experientially trained, c) the organization is using a more complex organizational structure, d) the type of internal stress present is not missing information or inaccurate information, and e) the level of internal stress is low. In addition, we find that the more complex the resource access structure the higher the organizational performance. Turnover actually degrades performance, thus if we consider all internal stressors other than turnover we find that information based stress tends to be more debilitating than individual based stress. In general, the correlation between accuracy and fallibility is .360 ($p < .005$). Organizations that make fewer mistakes also typically make fewer severe mistakes. However, the factors that lead to high accuracy do not also, necessarily, lead to low fallibility.

Rather, organizations that operate in an unbiased environment are more likely to make more correct decisions, but they are also more likely to make a major mistake when they make an error. Experientially trained organizations are both more accurate and less fallible. And so forth.

One issue with respect to the robustness of this analysis is the effect of task. In Table VII we have also summarized the effect of the factors we have been considering on performance when the organization is facing a binary task. At a gross level, the effects are comparable. However, there are two cases where the binary and trinary cases depart. These are, in a binary task environment the resource access structure and turnover appear to be less important than in the trinary task environment. These general results, however, do not mean that the top and bottom performing organizations are the same regardless of whether the task is binary or trinary. On the contrary, there is a limited interaction between the task and the organizational design. Future work should explore this relationship in greater detail.

III. Discussion

Working within the CORP framework allows the researcher to systematically explore the effects of a variety of factors on organizational performance. The numerous advantages of a computational approach need not be detailed here. What is important to note are the relative advantages of CORP. CORP employs a seemingly simple task and a set of highly stylized organizational designs. These designs are all defined from an information management point of view. Consequently, it is possible using CORP to trace through the effects of different constraints on information and the resultant impact on the organization. These constraints, such as the organizational structure, have recognizable analogs in real organizations. Indeed an advantage of comparing the relative performance of highly stylized designs over looking for the structure that generates the optimal performance is that these designs have corporate analogs; whereas, an arbitrary structure that has optimal behavior may not have such an analog.

In developing this model a number of simplifying assumptions have been made. These assumptions allowed us to lay bear the relation among various aspects of organizational design. However, they do limit the generalizability of the results. For example, we assumed a simple trinary choice task. While choice tasks are quite common in organizations, an important factor in those choices is that the number of choices is variable. Further, in many choice tasks there are a larger number of choices and part of the work involves elimination of non-contenders. This suggests that the observed results are most directly applicable to organizations engaged in fairly repetitive situations with only minimal options. Future work should consider whether the number of choices available and variability in the number of choices available affects the performance results. As another example we assumed a small number of personnel (nine analysts) and a small number of factors in the problem (nine). Most organizations have many more personnel divided across multiple ongoing problems that do not require cooperation. This suggests that the detailed results may be somewhat more applicable to project teams or crisis response organizations than to large organizations in which personnel are engaged in noncooperative ventures.

Individual (and hence organizational) performance is a function of the number of pieces of information (factors) that the individual must evaluate and the number of values each piece of information can take on (Carley, 1990). Thus organizational performance is affected by the ratio of analysts to problem factors, the ratio of analysts to mid-level managers, and the ratio of mid-level managers to the commander (top-level manager). Proportionally increasing these ratios while retaining the extant organizational design should have no impact on the results. As the number of pieces of information per person increases the task becomes more complex, the individual has a harder time learning, the potential for increase in errors due to missing information grows, and the chance of making a correct decision decreases. However, such changes should not affect the relative standing of various organizational designs. The small ratio of analysts to factors used in this paper suggests that the results are most applicable to simple tasks and or tasks where the work is relatively easy to compartmentalize.

There are many limits to this framework. For example, the task explored is time independent. The agents have, effectively, as much time as they desire to make their decisions. However, in most choice tasks, time matters. Organizations must respond in both a timely and an accurate fashion. Time pressure is a critical element of stress. Future work should consider how time constraints affect the performance of organizations with different designs. For example, we found that experiential organizations tended to outperform operational organizations in biased tasks and under external stress. Operational training, however, is expected to lead to faster performance than experiential training under stress. Such a speed advantage for operationally trained organizations may mitigate the inflexibility disadvantage that we observed. A second limitation is that within CORP the decision choices are represented numerically, e.g., with friendly as a 1 and hostile as a 3. While this representation seems natural it actually can affect the outcome. Procedures based on the "sum" of the values, as in this paper, tend to induce an internal bias in performance such that organizations perform better under "hostile" or "3" conditions than on average. Whether this is true for procedures based on the "number" of values of a certain type remains a topic for future research. Third, this paper assumes that the task is decomposable; i.e., that each piece of information is independent and equally weighted. Future research should consider more complex schemes where pieces of information interact.

IV. Conclusion

This research suggests that there is no one best organizational design. All aspects of design, the organizational structure, the resource access structure, and the training procedure all affect performance. In addition, factors which may be beyond the organization's control such as the type of stress and the task environment also affect performance. However, there is a certain amount of generality to the results. Further these results move us beyond organization theory by providing a

more detailed view of when what organizational design factors matter.

First, there is a training-task effect. In an unbiased task environment operational training is generally better than experiential; whereas, in a biased task experiential training is generally better. This is true regardless of whether the stress is internal or external. This is true whether one is considering accuracy or fallibility. Second segregated access is rarely valuable. The adage "divide and conquer" holds true only under highly specialized circumstances. In general, under stress situations, some information or resource sharing improves organizational performance and decreases fallibility. Third, in an unbiased task environment there is a tendency for simpler organizational structures (teams and hierarchies) to do better than complex structures (hierarchies and matrices); whereas, in the biased task environment more complex organizational structures, tend to do better. The main caveat on this is that for severe levels of different types of internal stress the exact structure that is the best (both organizational and resource) depends on what is going wrong, particularly if the agents are operationally trained. Fourth, organizational structures that are the most accurate are not necessarily the least fallible and vice-versa. For most organizations there is a tradeoff between making most of the decisions correctly and making the fewest severe mistakes. Organizations in choosing a design need to consider these tradeoffs. Organizations need to decide the type of risk they are comfortable with: do they need to be generally correct or do they need to never be completely wrong. Making this choice is intimately tied to the design choice. In this sense, this work is consistent with the findings of Roberts (1989, 1990) that high reliability organizations (few or no severe mistakes) need to be organized differently than other organizations [15], [16].

Researchers have suggested that increasing the number of levels in a hierarchy may be detrimental to the organization. As the number of levels increases the degree of information distortion [12], condensation [8], or uncertainty absorption [13] should increase and the organization's accuracy should decrease and the chance of fallibility increase. In other words, teams should outperform hierarchies. In the CORP model, the greater the number of hierarchical levels the greater the information distortion, condensation, and uncertainty absorption. Earlier work using CORP and experiential agents facing a binary task showed that under non-stress conditions indeed teams did outperform hierarchies [3], [4]. In this paper, we again examine hierarchies with either 2 or 3 levels under stress. However, we do not find a consistent effect due to hierarchical levels. Moreover, the dominant result is that increasing the number of levels leads to improved performance. A question for future research is whether there is a point at which the degradation in performance, due information distortion etc. as the number of levels in the organization is increased, overcomes the performance advantage, due to the embedded learning that accrues when these levels are increased.

In the final analysis, the results of this model should be related to empirical data. As a preliminary analysis we located archival information (newspaper reports and scholarly articles) on ten organizations subject to a technological "crisis" condition (such as the Exxon Valdez). These organizations tend to be simultaneously undergoing both internal and external stress. Typical internal stresses include missing information, incorrect information, or missing personnel. Each organization could be characterized by whether the personnel were directed to act on the basis of their experience or were expected to following standard operating procedures. For each organization we coded their organizational design using the categories discussed in this paper. Performance was measured in term of the socio-economic impact of the crisis and the chance for avoidance. Most of these organizations are essentially operating in a biased environment; i.e., certain outcomes are much more likely than other outcomes.

The simulation results presented in this paper can be thought of as a reasonable point estimate of organizational performance in a specific situation. Based on the foregoing analysis we would predict that more complex structures such as hierarchies or matrices with blocked or distributed would be best under high stress conditions, and that experiential training would be best. We find better performance among the experientially trained organizations than the operational and we find the more complex structures outperforming the simpler. Clearly, this analysis is preliminary and future research should be brought to bear on this issue. However, the results are encouraging and suggest the plausibility of the CORP framework.

REFERENCES

[1] A. H. Bond and L. Gasser, *Readings in Distributed Intelligence*, San Maeto, CA: Morgan Kaufmann Publishers, 1988.

[2] K. Carley, "Coordinating for Success: Trading Information Redundancy for Task Simplicity," in *Proceedings of the 23rd Annual Hawaii International Conference on System Sciences*, 1990.

[3] ———, "Designing Organizational Structures to Cope with Communication Breakdowns: A Simulation Model," *Industrial Crisis Quarterly*, vol. 5, pp. 19–57, 1991.

[4] ———, "Organizational Learning and Personnel Turnover," *Organization Science*, vol. 3, no. 1, pp. 20–46, 1992.

[5] K. Carley and Z. Lin, "Maydays and Murphies: A Study of the Effect of Organizational Design, Task and Stress on Organizational Performance," *Working Paper*, Department of Social and Decision Sciences, Carnegie Mellon University, 1993.

[6] M. D. Cohen and J. G. March, *Leadership and Ambiguity, New York*, NY: McGraw Hill, 1974.

[7] M. D. Cohen, J. G. March and J. P. Olsen, "A Garbage Can Model of Organizational Choice," *Administrative Sciences Quarterly*, vol. 17. no. 1, pp. 1–25, 1972.

[8] A. Downs, *Inside Bureaucracy*, Boston, MA: Little Brown, 1967.

[9] J. Galbraith, *Designing Complex Organizations*. Reading, MA: Addison-Wesley, 1973.

[10] J. R. Hollenbeck, D. J. Sego, D. R. Ilgen and D. A. Major, "Team Interactive Decision Making Exercise for Teams Incorporating Distributed Expertise (TIDE^2): A Program and Paradigm for Team Research," *Tech. Rep.*, no. 91–1, East Lansing: Michigan State University, Department of Management and Psychology, 1991.

[11] D. R. Ilgen, D. A. Major, J. R. Hollenbeck and D. J. Sego, "Decision Making in Teams: Raising and Individual Decision Making Model to Team Level," *Tech. Rep.*, no. 91–2, East Lansing: Michigan State University, Department of Management and Psychology, 1991.

[12] F. M. Jablin, L. L. Putnam, K. H. Roberts and L. W. Porter (Eds.), *Handbook of Organizational Communication: An Interdisciplinary Perspective*, Sage, Beverly Hills, CA, 1986.

[13] J. G. March and H. A. Simon, *Organizations*, Wiley, 1958.

[14] C. Perrow, *Normal Accidents: Living with High Risk Technologies*, New York, NY: Basic Books, Inc., 1984.

[15] K. Roberts, "New Challenges to Organizational Research: High Reliability Organizations" *Industrial Crisis Quarterly*, vol. 3. no. 3, pp. 111–125, 1989.

[16] ——, "Some Characteristics of One Type of High Reliability Organizations," *Organization Science*, vol. 1. no. 2, pp. 160–176, 1990.
[17] Z. Tang, K. R. Pattipati and D. L. Kleinman, "A Distributed M-ary Hypothesis Testing problem with Correlated Observations," *IEEE Trans. on Automatic Control*, pp. 632–636, Jun. 1992.
[18] J. Thompson, *Organizations in Action*, New York, NY: McGraw-Hill, 1967.

Kathleen M. Carley was born in Pueblo, CO, in 1956. She received two S.B. degrees in 1978 from the Massachusetts Institute of Technology, Boston, MA, one in political science and one in economics. She received her Ph.D. in sociology from Harvard University in 1984. She is now Associate Professor of Sociology and Information and Decision Systems at Carnegie Mellon University, Pittsburgh, PA.

Her research interests focus on the joint cognitive and structural basis for social and organizational behavior in dynamic settings. Research interests include: (1) Examining the impact of organizational design and the member's skills on organizational performance and teamwork under normal and crisis conditions using computational models (see e.g., 1992 article in Organization Science and the book Carley and Prietul a (eds) 1994 Computational Organization Theory). (2) Examining the interplay of social structure, cognition, and technology in a ffecting information diffusion and group formation (see e.g., the book Kaufer and Carley, 1993, Communication at a Distance). (3) Encoding mental models from texts (see e.g., Sociological Methods and Resea rch, 1988, Social Forces 1992). (4) Development of statistical techniques for network analysis (see e.g., Banks and Carley, 1994, Journal of Classification).

Zhiang Lin was born in Shanghai, People's Republic of China, in 1963. He received a B.S. in Computer Science in 1986 and an M.S. in Human and Quantitative Geography in 1988, both from East China Normal University, Shanghai. After working at Shanghai Wang Computer Co. Ltd. for one year, he went to Carnegie Mellon University on a full scholarship, where he earned a Master of Philosophy degree in Public Policy and Management in 1991 and successfully defended his Ph.D. thesis in Organizations and Policy Analysis He now is an Assistant Professor of Organizations at The Hong Kong University of Science and Technology.

His research interests focus on organizational theory and the modeling of organizations. He has developed two different simulation testbeds, CORP and DYCORP, for examining the interplay of organizational design, agent intelligence, training, and task on organizational performance. He has examined the impact of agent style (1993, International Journal of Intelligent Systems in Accounting, Finance and Management), and the predictive power of various measures of organizational design (1994, "A Theoretical Evaluation of Measures of Organizational Design" in Carley and Prietula (eds) Computational Organization Theory).