Industrial Crisis Quarterly 5 (1991) 19–57
Elsevier

Designing organizational structures to cope with communication breakdowns: a simulation model

Kathleen M. Carley
Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

Abstract

Carley, K.M., 1991. Designing organizational structures to cope with communication breakdowns: a simulation model. Industrial Crisis Quarterly, 5: 19–57.

During crises, organizations commonly face communication breakdown and erroneous information. For example, equipment malfunctions, decision makers become unavailable, and so on. Such breakdowns should degrade the organization's ability to make decisions. However, the design of the organization – who can communicate with whom and who has access to what information – should determine the extent to which organizational performance is degraded by communication breakdowns. In this paper a model of organizational decision making when the organization is engaged in tasks so complex that the separate decisions of multiple decision makers must be integrated to locate the organization's decision is developed. Using simulation, the implications of this model for how organizations should be designed to best cope with crisis are examined.

Introduction

At some point or another, and certainly more frequently than desired, organizations face periods of crisis. Some of these crises are associated with actual terrible accidents, such as the space shuttle *Challenger* accident or disasters such as the Mexico City earthquake. Other crises are associated with potential,

This research was supported by the NSF under grant No. SES-8707005. An earlier version of this paper was presented at the Second International Conference on Industrial and Organizational Crisis Management, New York University, 3–4 November 1989.

rather than realized, catastrophic consequences, such as the 2,159 missile display conferences which were called at NORAD during the first six months of 1980. Still other crises are associated with problems of a more limited scope, such as when university computers go down the week before term papers are due. In all cases, analysts often point to the fact that during these crisis periods the performance of the organization is compromised by communication breakdowns and poor information. In other words, bad decisions are made because the right information is not in the right place at the right time.

This paper examines the following question. How should an organization be designed so that it is the least affected by crisis? This question is addressed analytically using a model of organizational decision making that makes it possible to contrast the likelihood that the right information gets to the right place at the right time despite various types of communication breakdowns in organizations with different structures. First, however, in order to gain some perspective on the role of communication breakdowns during periods of crisis, two examples are considered: (1) a false alarm about a missile attack, and (2) the *Challenger* accident. These cases suggest what features the model must have if it is going to provide even limited insight into the question of organizational design.

NORAD (North American Aerospace Defense Command) is an early warning command center that analyzes and responds to incoming signals that indicate that the US may be under military attack. In a report to the Senate Committee on Armed Services it was noted that missile display conferences were called 2,159 times during the first six months of 1980 in response to potential threats. Let us consider one such case that occurred at 2:26 AM on 3 June 1980, as reported by Perrow:

> Strategic Air Command received an indication that two submarine-launched missiles were on the way; it came from NORAD headquarters, SAC called to confirm, and NORAD was unable to confirm, even though the message came from their computers. The SAC duty officer took the precaution of ordering all B-52 alert crews to board their airplaines and start their engines ... Shortly thereafter the SAC display indicated no missiles, and no other parts of the warning systems indicated any either, so the B-52 crews were ordered to shut down their engines.
>
> A few minutes later the SAC display monitor, receiving messages from NORAD headquarters, indicated Soviet land-based missiles on their way to targets in the United States, and a little later the monitors in the Pentagon showed submarines missiles coming in ... The duty officer in the Pentagon convened a missile display conference, the lowest level of alert involving a conference telephone call, and then went on to the next level, a missile threat assessment conference bringing in more officers. The commander of NORAD said that there was in fact no threat, and one minute later the SAC alert was terminated ...
>
> The cause of the false alarm could not be determined. But NORAD knew it was a false alarm because neither the satellite nor the radar picked up signals of land or submarine-based missiles being launched or penetrating our perimeters.
>
> (Perrow, 1984, pp. 286–287)

This false alarm was caused by a defective computer chip in the multiplexor sending incorrect data to some, but not all, command posts. In other words, the period of crisis was characterized by the presence of incorrect incoming information. Such incorrect information, in this example, was not particularly dangerous due to the fact that other sources of incoming information suggested a totally different scenario (no missiles).

In the second example, the *Challenger* accident, at least two different types of communication breakdowns occurred – intermittent participation and incomplete information. As noted in the *Report of the Presidential Commission on the Space Shuttle Challenger Accident*:

> ... testimony reveals failures in communication that resulted in a decision to launch 51-L based on incomplete and sometimes misleading information, a conflict between engineering data and management judgments, and a NASA management structure that permitted internal flight safety problems to bypass key Shuttle managers.

And in particular:

> ... in the launch preparations for 51-L relevant concerns of Level III NASA personnel and element contractors were not, in the following crucial areas, adequately communicated to the NASA Level I and II management responsible for the launch:
> - The objections to launch voiced by Morton Thiokol engineers about the detrimental effects of cold temperatures on the performance of the Solid Rocket Motor joint seal.
> - The degree of concern of Thiokol and Marshall about the erosion of the joint seals in prior Shuttle flights, notably 51-C (January, 1985) and 51-B (April, 1985).

Further examination of the events preceding the accident indicates a series of decision points at which the flow of participants was continually changing. For example, Allan McDonald, Thiokol's liaison for the Solid Rocket Booster project at Kennedy Space Center, after being informed of potential problems due to the low temperatures, tried to contact Lawrence Mulloy, the Solid Rocket Booster project manager at Marshall, only to find that he was unavailable. Then the first phase of the 5:45 EST teleconference began without McDonald, who did not arrive at Kennedy until 8:00 PM. Later portions of conferences were held offline, thus breaking communication channels. Then the flight readiness rationale was sent at 11:45 after an offline caucus at Thiokol. And, as noted by Mulloy, "I had no knowledge of what interchange occurred during the caucus at Thiokol, because all sites were on mute. We were on mute at KSC. No communications occurred between myself and Mr Hardy at Huntsville, nor did any communication occur between KSC and Thiokol during that conference" (Report of the Presidential Commission, 1986, p. 96). In summary, communication breakdowns were quite prevalent during the window in which the decision to launch the shuttle was made.

These two examples, and more detailed case studies such as those of Bhopal (Shrivastava, 1987), Three Mile Island (Kemeny, 1981; Sills et al., 1982), and

Grenada (Metcalf, 1986), provide detailed insights into the factors that determine organizational performance during crises. Such studies demonstrate that a variety of structural, systemic, political and psychological factors determine organizational performance. In particular, these studies suggest that while technological and political factors often determine whether a crisis occurs, the ability of the organization to mitigate the crisis, operate effectively during the crisis, or even avoid the crisis, is dependent on factors such as the organizational structure (Metcalf, 1986; Report of the Presidential Commission, 1986), the cognitive capabilities and experience of the individual decision makers (Metcalf, 1986; Shrivastava, 1987), the frequency, length and form of the communication breakdowns during the crisis (Perrow ,1984; Report of the Presidential Commission, 1986; Shrivastava, 1987), the quality of the incoming information (Kemeny, 1981; Perrow, 1984; Metcalf, 1986; Report of the Presidential Commission, 1986), and the degree to which different personnel in the organization share the same information (Kemeny, 1981; Metcalf, 1986; Report of the Presidential Commission, 1986; Shrivastava, 1987).

Further, these studies clearly demonstrate that during the crisis period, organizational performance is compromised by communication breakdowns and incorrect information. In order to reduce the impact of communication breakdowns, decrease response time and increase organizational performance during crisis, organizations are expected to have emergency preparedness or contingency plans. According to Michael (1986), such a plan should include, among other things, procedures for notifying and communicating with management and the establishment of a command and coordination structure and an allocation of responsibility for the various tasks. In other words, the plan should specify an organizational design. Yet studies of organizational design provide little insight into what organizational structure is most effective during crisis.

Case studies, like those previously discussed, do not provide us with a systemic account of the relative impact of each design factor separately or collectively on organizational performance. In order to move beyond case studies, organizational theorists have been turning to the development and analysis of mathematical models of organizations in order to systematically examine the relationships between various organizational features and performance (Cohen et al., 1972; Padgett, 1980a; Anderson and Fischer, 1986; Carley, 1986, 1990; Masuch and LaPotin, 1989). Such studies illustrate that organizations with different designs have different performance characteristics. For example, Padgett found that within hierarchies the chief executive officer (CEO) can most effectively get the organization to make decisions that meet the CEO's goals by exercising a hands-off policy and hiring liberal assistants for low saliency programs and conservative assistants for high saliency programs. And

Carley (1989) found that teams learn faster and better than hierarchies but are less resilient in the face of turnover.

The question of the best organizational design to cope with the communication breakdowns that occur during crisis has not been addressed analytically. However, much of the previous work does have implications for performance during crisis. For example, the work by Cohen et al. (1972) suggests that regardless of the organizational design, when the information load increases (as it does during a crisis), fewer problems are resolved. And a study by Carley (1986) suggests that small, differentiated organizations exhibit the highest efficiency during crisis. In both of these studies, the organizations examined where fully functional; that is, there were no communication breakdowns.

To summarize, case studies of crises provide a wealth of information concerning organizational operation but they do not admit systematic investigation of the relative impact of different organizational designs on performance and the ability of the organization to make decisions. Moreover, such studies suggest a model of organizational behavior that is at odds with the models extant in the organizational decision making literature. Specifically, these cases suggest that: (1) organizations make decisions by integrating the decisions and partial information of multiple decision makers, not through reaching consensus and sharing all information; (2) institutional memory is distributed, and consequently, what the organization knows is a function of both which individual decision maker knows what and the relationships (such as communication and command channels) among the decision makers; and (3) personnel are engaged in tasks for which they have been trained and do somewhat repetitively but with minor variations. In contrast, most studies of organizational and group decision making focus on one-of-a-kind tasks or completely repetitive tasks (Marschak, 1955; Wilensky, 1967) or on tasks where consensus must be reached before a decision is made (Bavelas, 1950; Cohen et al., 1962; Cohen et al., 1969; DeGroot, 1974; Hastie, 1986). Consequently, a model of organizational behavior that more closely parallels the observed behavior is called for.

In this paper, a model of organizational performance is presented and used to systematically examine how organizations should be designed to cope with the communication breakdowns that occur during crises. Performance characteristics during periods of crisis and periods of normal operating conditions are contrasted in order to determine the degree to which crises degrade performance. More specifically, the impact of communication breakdowns on organizational performance for organizations with either a centralized hierarchical, a team, or a dual-command hierarchical structure are examined using a model of organizational decision making that is based on a model of experiential-based individual decision making.

The proposed model is used to examine the impact of communication breakdowns on organizations engaged in tasks where they face a sequence of highly similar, but rarely identical, problems that are so complex that no one decision maker in the organization may be able to cope with, let alone have access to, all of the information needed to make a decision. I refer to sets of problems with these characteristics as "quasi-repetitive integrated decision making tasks." In a quasi-repetitive integrated decision making task it is not necessary (and often not possible) for individuals to reach consensus, and information pooling is done by institutional design (Panning, 1986) rather than by individual choice of communication partner. In the quasi-repetitive integrated decision making tasks, the same basic type of problem is faced over and over again but some of the information, constraints, parameters, etc. are different each decision period, thus producing slightly different decisions. *Quasi-repetitiveness* can be thought of as a continuous scale bounded on one end by repetitive tasks and on the other by non-repetitive tasks. A task would be repetitive if exactly the same problem was faced over and over again, whereas a task is non-repetitive if the problem is unique; some crises by their very nature may be on the less repetitive end of the spectrum. A task is said to be *integrated* if the final organizational decision is determined by somehow integrating into a single decision a plethora of previous smaller or component decisions made by various decision makers within the organization. Many organizations face this type of task – fire companies, air traffic control teams, design teams, NORAD, NASA, and so forth. The proposed model is an extension of an earlier model in which organizations are viewed as adaptive systems whose performance is dependent on the integrated performance of the individual decision makers in the organization who themselves learn from experience (Carley, 1990, forthcoming). This model combines an individual or cognitive information processing perspective (Simon, 1947; March and Simon, 1958; Cyert and March, 1963; Cohen et al., 1972) with an organization structural perspective (Weber, 1922; Galbraith, 1973) and a distributed decision making perspective (Bond and Gasser, 1988) in order to extend the argument that performance is a function of organizational design (Kimberly, 1987; Hinnings and Greenwood, 1988) to performance during crisis is a function of design.

There have been a few studies of tasks with characteristics like those examined in this paper; examples include work on air traffic control (Steeb et al., 1980; Thorndyke et al., 1981; LaPorte and Consolini, 1988), sensor data interpretation (Smith, 1980; Carley et al., 1988), and planning and budgeting (Padgett, 1980b; Metcalf, 1986). Such studies suggest that the design of the organization or its coordination scheme can dramatically affect organizational performance. Despite the prevalence of such tasks in organizations, the ability of organizations confronted with such tasks to cope with communication breakdowns has not been explored analytically. Yet when communication breakdowns occur in organizations which face quasi-repetitive integrated decision making tasks the results are often catastrophic. Reconsider the two examples. The personnel at NORAD and NASA are engaged in extremely complex integrated tasks; for example, the decision to launch the shuttle is based on the decisions of multiple individuals in different groups that their portion of the system is ready. In both examples, the personnel are doing what they have been trained to do; shuttle launches and missile display conferences are quasi-repetitive events where the particular anomalies of the crisis period made the specific problem faced by the organization just a little different from the problems it had previously experienced – the weather was a little colder, a particular computer chip failed.

Communication breakdown model

Organizations are viewed from a distributed decision making framework; that is, organizations are engaged in a quasi-repetitive integrated decision making task. Organizational life is divided into a sequence of decision periods. In each decision period the organization faces a new problem that is similar, but not identical, to previous problems. During each period, information related to the new problem is evaluated, a decision is made, and the members of the organization are informed of the correctness of their decision for that problem. This operational procedure represents the normal operating conditions for the organization and is followed during non-crisis periods. Crises disrupt this process. Three types of crises are considered: incorrect information, decision makers become unavailable, and communication channels become unavailable.

Task

Each time period, the organization as a whole needs to make a decision. This decision is modeled as a simple yes/no choice such as "Is the radar bleep a missile?" or "Should the shuttle be launched?" In making this decision, the organization has access to a set of information. This information can be thought of as the presence or absence of a set of factors such as "Is the radar bleep moving fast or slow?" or "Is the temperature greater than 40 degrees?" If enough factors are present, a positive or *yes* decision is made, otherwise a *no* decision is made. Mathematically, the information for this task can be represented as a binary string such that each "bit" represents whether a particular factor is present or not. If the factor is present, e.g. if the radar bleep is moving fast, the bit is a 1; otherwise it is a 0. The particular configuration of 1's and 0's is the information the organization needs to decide what to do. In this paper, all tasks examined have a task complexity of 27, that is, the number of bytes in the string is 27[1]. Thus, the number of different problems that the organization might see is 2^{27}. Similarly, the organization's decision (e.g. the bleep is a hostile missile) can be mathematically represented as a binary choice – yes (1) or no (0).

Each decision period, the organization is faced with a specific problem; for instance, can NORAD confirm the indication that at 2:26 AM on 3 June 1980 two submarine-launched missiles are on their way? This specific problem is represented as a specific pattern of 1's and 0's. In order to solve this problem a vast array of information must be analyzed. There is so much information that multiple decision makers are required. Indeed, each problem can be divided into a set of sub-problems, such as what the Pentagon monitors say, what the satellites say, and so on. Such a division is represented in this model by giving each analyst only a portion of the world, that is, a sub-problem. The size of sub-problems and the degree of overlap in who knows what depends on the level of redundancy as defined by the information access structure.

Decision makers are limited and can act only on the basis of the information that they know (Simon, 1947; March and Simon, 1958; Cyert and March, 1963). This is represented by having analysts make decisions – yes or no – on the basis of just the information in the sub-problem that they see. Decisions and not information are communicated to the managers; that is, each decision maker makes a recommendation for what he or she thinks the final decision should be. The individual decision maker by passing on a yes or no decision rather than the number of 1's or 0's observed has compressed information; hence there is information loss. The more information the individual needs to process, the more complex the task per individual, the greater the information loss and the lower the overall organizational performance (Carley, forthcoming).

Because information is distributed, the organization must integrate the decisions made by all of its members in order to reach its ultimate decision. I refer to the integrated decision made by the organization as its *final decision*. The final decision provided by the organization is, within a hierarchy, the decision made by the CEO and is a "1" if the CEO decides that the answer is a 1, and a "0" if the CEO decides the answer is a 0. For the team, the final decision is the majority vote made by the analysis – "1" if more analysts think the answer is a 1, and "0" if more analysts think the answer is a 0. Due to information

[1]In determining what level of task complexity to examine, two requirements were taken into account. First, in order to guarantee that there is a correct decision, the number of bits in the full problem must be odd. Second, it must be possible to divide the problem such that, with the same number of analysts, different levels of information redundancy can be explored given that there are nine analysts. When there is no redundancy and nine analysts, the possible choices of task complexity are odd multiples of nine (9, 27, 45, 63 ...). When redundancy is admitted, the task complexity must be greater than nine. Thus, 27 is the simplest case that meets these criteria.

loss, the final decision may not be correct. The decision that the organization would have made if there was no information loss is referred to as the *correct decision.* Performance of the organization is characterized in terms of what fraction of the final decisions are correct.

Organizations do not blindly, and rationally, evaluate new information in order to reach a decision. Rather, organizational behavior is historically based (Lindblom, 1959; Steinbruner, 1974; Levitt and March, 1988). In making decisions, organizations rely on experience, incrementally adapting their response to similar problems as they receive feedback on their previous decisions. In this model, the feedback process is represented by informing each decision maker each decision period as to what was the correct decision. All members of the organization find out what the organization should have done; for example, they are informed that the bleeps are not missiles, or the shuttle should not have been launched. As will be discussed, the decision makers learn, that is, they gain experience as they make decisions and receive such feedback.

Organizational design

Each organization has a particular design or coordination scheme as defined by the C3I structure, associated order of processing, and rules for processing information and making decisions. In this paper, the design is characterized by the organizational structure (who commands/communicates to whom and the procedure for producing a final organizational decision) and the information access structure (who has access to what information and hence the level of information redundancy).

Organizational structure

Organizations vary in terms of their formal structure for who commands or communicates to whom. In this paper, three organizational structures are examined: the centralized hierarchy, the dual-command hierarchy and the team (see Fig. 1). These structures are not meant to exhaust the set of potential or actual organizational structures. Rather, they represent idealized structural types that are interesting due to their prevalence in real organizations, and that may be differentially affected by communication breakdowns. In addition, both the hierarchy (Weber, 1922; Burns and Stalker, 1961; Padgett, 1980a; Carley, 1986; Malone, 1986, 1987) and the team (Marschak, 1955; Strand, 1971; Cohen et al., 1972; Bar-Shalom and Tse, 1973; Gloves and Ledyard, 1977; Arrow and Radner, 1979; Tsitsiklis and Athans, 1984; Radner, 1987) have been extensively studied, but their performance has rarely been contrasted. The dual-command hierarchy, such as that employed in the material acquisition process by the army, has rarely been studied, although in certain ways it is similar to

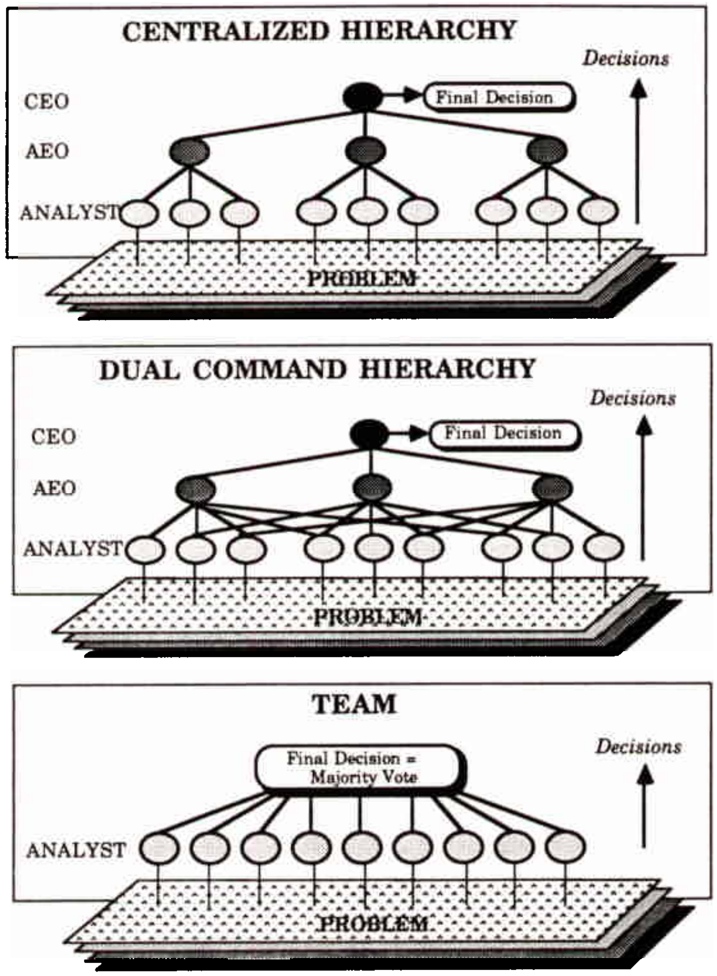

Fig. 1. *Organizational structures.* Three organizational structures are examined – the centralized hierarchy (top), the dual-command hierarchy (middle) and the distributed team (bottom). Analysts are represented as lightly shaded circles, AEOs as darkly shaded circles, and the CEO as a black circle. In both types of organization each analyst sees a portion of the problem and makes a decision. In the centralized hierarchy and the dual-command hierarchy this decision is passed to the AEO who then makes a decision which is passed to the CEO, who makes the final decision for the organization. In the team, the analyst's decision is his "vote" as to what the final decision should be, majority rules. In all cases, only the analysts have access to the 'raw data" associated with the problem.

the matrix structure (Davis and Lawrence, 1977). These three structures represent different points on the organizational spectrum in terms of the degree to which institutional memory and command is centralized.

The primary difference between the team and the hierarchical structures (centralized hierachy and dual-command hierarchy) is the presence of upper level management in the hierarchies and the absence of such management in the team. Such upper level management may potentially reduce the impact of communication breakdowns as experienced management has a historical image of the entire structure and managers integrate lower level decisions. The primary difference between the centralized hierarchy and the dual-command hierarchy is the presence of multiple paths by which information can reach management in the dual-command hierarchy and only a single such path in the centralized hierarchy. The presence of multiple paths may potentially reduce the impact of communication breakdowns by reducing the likelihood that some information will be completely lost.

Centralized hierarchy: The centralized hierarchy is modeled as a three-tier organization composed of a CEO, a set of assistant executive officers (AEOs) and a set of analysts. Each analyst in each decision period receives information (a sub-problem), makes a decision (yes or no), and sends this decision to his or her AEO. The AEO takes the analysts' decisions, makes an integrated decision (yes or no), and sends this decision to the CEO. The CEO takes the AEOs' decisions and makes the final integrated decision (yes or no). In this paper, the specific centralized hierarchy examined has thirteen decision makers, with three under each manager or executive, as in Fig. 1. There are nine analysts. A hierarchy of thirteen decision makers is the minimum non-trivial hierarchy that can be examined such that the hierarchy has three levels and an odd number of decision makers under each manager.

Dual-command hierarchy: The dual-command hierachy, like the centralized hierarchy, is modeled as a three-tier organization composed of a CEO, a set of AEOs, and a set of analysts. Each analyst in each decision period receives information (a sub-problem), makes a decision (yes or no), and sends this decision to his or her AEO. In addition some analysts also send their decisions to another AEO who is listed as interested in the project. The AEO takes the analysts' decisions, makes an integrated decision (yes or no), and sends this decision to the CEO. The CEO takes the AEOs' decisions and makes the final integrated decision (yes or no). In this paper, the specific dual-command hierarchy examined has thirteen decision makers, with three under each manager or executive, as in Fig. 1. There are nine analysts. For each AEO, one analyst reports only to him or her, and the other two analysts each also report to one other AEO. This size of organization is chosen so that it matches the hierarchy.

Team: The team is modeled as a single-tier organization composed of a set of analysts. In each decision period, each analyst receives information (a sub-problem), and makes a decision (yes or no) independent of the other analysts. The organization's decision (the final decision) is the majority vote of the analysts. In this paper, the specific team analyzed has nine decision makers, all of whom are analysts, as in Fig. 1. The team has nine decision makers in order to match the number of analysts in the hierarchy. The number of analysts, rather than total decision makers, is matched, so that the complexity of the sub-problem seen by each analyst is identical for the team and hierarchy for the same size problem while maintaining no overlap in the sub-problems seen by the different analysts.

Information access structure
When problems are divided across the personnel in the organization some people may share information. For example, engineers at both Thiokol and Marshall knew about the erosion of the joint seals in prior shuttle flights. Which personnel know which pieces of information, and hence the level of redundancy, is defined by the information access structure. Organizations differ in these structures. In some organizations information is sparsely distributed such that there is little or no overlap in who knows what. These are segregated structures where each decision maker is specialized (Cohen et al., 1972), whereas in other organizations there is greater overlap in who knows what. In this paper three distinct information access structures are examined: (1) segregated, (2) blocked, and (3) distributed. These are displayed in Fig. 2. Each structure determines which analyst sees which portion of the problem and how much of the problem. Like the organizational structures, these information access structures are not meant to exhaust the set of potential or actual information access structures. Rather, they represent idealized types that are interesting due to their prevalence in real organizations and that may be differentially affected by communication breakdowns. Of the three, the segregated access structure has been the most frequently studied and is often referred to as a specialized structure. This structure has been subjected to both analytical (Cohen et al., 1972; Carley, 1986) and experimental analysis (Cohen, 1962; Shaw, 1981) and is typically contrasted with semi-specialized or complete (where everyone knows everything) access structures. The results, however, are ambiguous. While Cohen et al. (1972) find that specialization leads to more problems being resolved, Carley et al. (1988) find that specialization leads to mistakes. Case studies point out that there is often redundancy in who knows what in the organization such as one observes in both the blocked and the distributed structures.

Segregated: In the segregated structure, each analyst has access to a unique set of information; that is, there is no redundancy. In this case, should anything

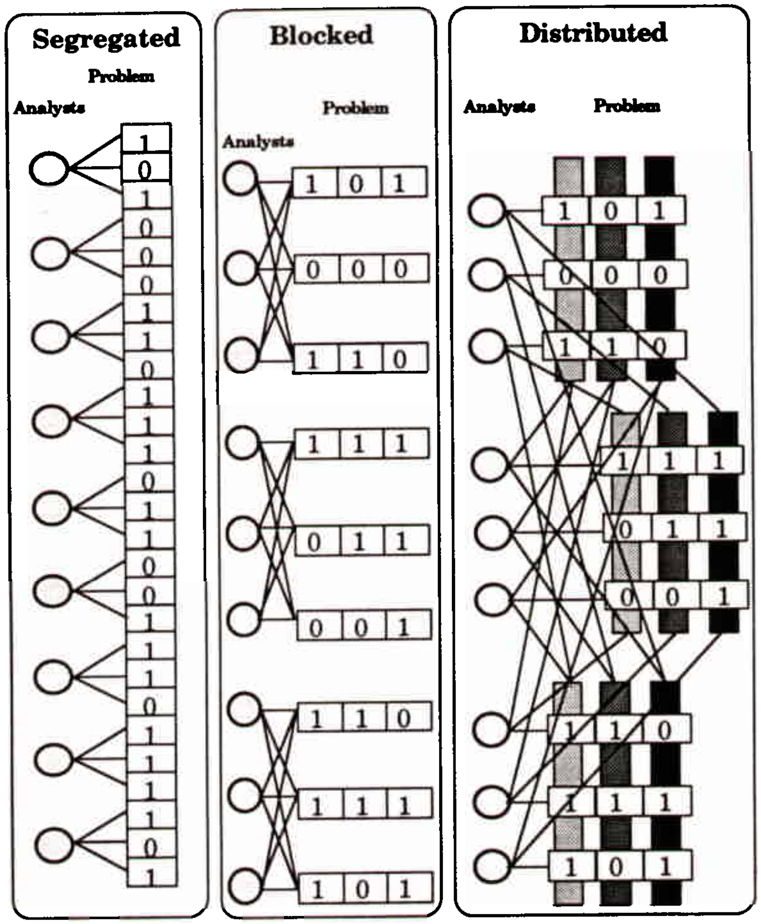

Fig. 2. *Information access structure.* Three information access structures are examined – the segregated (left), blocked (middle) and the distributed (right). In the segregated structure, each analyst has access to only three pieces of information and no information is shared by two or more analysts. In the blocked structure, each analyst has access to nine pieces of information and three analysts have access to identical information. In the distributed structure, each analyst has access to nine pieces of informationand no two analysts have access to identical information.

compromise an analyst, such as a communication breakdown or incorrect information, the organization as a whole loses information and so may not be able to make the correct decision. This structure is implemented by giving each analyst access to a unique set of three bits. Since the task is composed of 27 bits and there are nine analysts regardless of organizational structure, giving three bits to each analyst results in no redundancy.

Blocked: In the blocked structure, all members of an organizational division or project group share identical information. In this case, should anything compromise an analyst, the organization does not lose information and there are two other analysts who also report to management or vote in exactly the same way. This structure is implemented by giving each analyst access to nine bits of information and three analysts have access to the identical set of information. In the hierarchies, all analysts under one AEO have exactly the same information.

Distributed: In the distributed structure, each piece of information is known by multiple analysts but each of those analysts is in a different organizational division or project group. As in the blocked structure, each piece of information is known by three analysts; but, unlike the blocked structure, no two analysts (let alone three) have access to exactly the same information. Consequently, two analysts who share some information may still make different decisions as they also know different information. In this case, should anything compromise an analyst, the organization does not lose information but it might lose some of the value of that information. This structure is implemented by giving each analyst access to nine bits of information such that each bit is also known by two other analysts.

When the information access structure is blocked and distributed the analysts have more information to process. Similarly, in the dual-command hierarchy, the AEOs are faced with additional information-processing demands over their counterparts in the centralized hierarchy. In both cases, in Perrow's terms (1972), there is low *analyzability* and heightened *task uncertainty* until the organization is fully trained. However, once fully trained, such additional information actually heightens certainty. Due to these processing differences these organizations are expected to be differentially affected by crisis.

Decision procedure and learning

As previously noted, individuals in the organization learn from their experiences. Further, the organization's performance is dependent on the experience and capabilities of the individual members (see Shaw, 1981, and Hastie, 1986, for reviews). Thus the ability of the organization to learn depends, at least partially, on the memories of individual decision makers within the organization and their ability to learn (Hastie et al., 1984; Johnson and Hasher, 1987). All personnel, regardless of their job or position within the organization (analyst, AEO or CEO), learn from experience and can be characterized as intendedly adaptive (March and Olsen, 1975).

This learning process can be characterized as follows. The individual decision maker is faced with the problem of deciding whether the object on the radar is a missile, given the information that the object is moving fast, is large, and appears paired with another object. The decision maker categorizes this information; e.g this reminds me of another time when ... Then the decision maker uses information from previous similar cases to interpret the new information and make a decision. For example, the decision maker might reason that last time this happened it was a computer malfunction, and therefore this time it may also be a computer malfunction. When the decision maker receives feedback that the current problem is also a computer malfunction, this line of reasoning is reinforced. Work in behavioral decision theory suggests that individuals faced with such problems are imperfect statisticians insensitive to sample size (Tversky and Kahneman, 1971), but adjusting their expectations on the basis of additional information (Tversky and Kahneman, 1974), and effectively overconfident of their ability to correctly predict outcomes (Lichtenstein and Fischhoff, 1977).

In this paper, this type of learning process is modeled by letting each decision maker keep a cumulative record of the sub-problems that it receives, its decisions, and the true answer. For each decision maker, each sub-problem falls into a particular class. A class is a particular pattern of 1's and 0's, such as 010. The complexity of the sub-problem seen by the decision maker is defined by the size of the sub-problem (that is, number of bits of information he or she must analyze). The complexity of the sub-problem faced by the decision maker depends on his or her position in the organizational structure and the coordination scheme used by the organization (see Table 1). For example, in a cen-

Table 1

Complexity of sub-problem faced by various decision makers

Position	Access structure		
	Segregated	Blocked	Distributed
Centralized hierarchy			
Analyst	3	9	9
AEO	3	3	3
CEO	3	3	3
Dual-command hierarchy			
Analyst	3	9	9
AEO	5	5	5
CEO	3	3	3
Team			
Analyst	3	9	9

tralized hierarchy with a segregated access structure each decision maker sees the same number of positions or bits, three, and so have $2^3 = 8$ classes of sub-problems.

As the decision maker encounters sub-problems, it builds up, for each class of sub-problems, an expectation as to whether the true decision when it sees a problem in that class is a 0 or a 1. The expectation that the answer is a 0 is defined as the proportion of times in this decision maker's experience that, given this class of problems, the correct decision was a 0. The expectation that the answer is a 1 is defined as the proportion of times in this decision maker's experience that, given this class of problems, the correct decision was a 1. When the decision maker is faced with a sub-problem, the decision maker uses this experiential information to make a decision using the following procedure:

(1) Determine what class the problem is in.

(2) If the expectation of a 0 is greater than the expectation of a 1, return 0 as the decision.

(3) If the expectation of a 0 is less than the expectation of a 1, return 1 as the decision.

(4) If the expectation of a 0 is equal to the expectation of a 1, return either a 0 or a 1 as the decision with equal likelihood.

When modeled this way the decision makers exhibit the behavioral characteristics identified earlier. Further, by using this decision and feedback procedure, the decision makers are learning to match incoming information to possible decisions in much the same way that parallel distributed processing systems learn to match particular patterns to particular outputs (Rumelhart et al., 1986). The proposed procedure is in effect weighting each input separately. As such, the learning procedure defined guarantees that the decision maker will come to attend more to that incoming information that will match the correct response.

Even though the organization and its members know what kind of decision is supposed to be made, they may not know when to make the decision. For example, they may know that they are to determine whether the missile is hostile but they may not know what configuration of characteristics correspond to the missile being hostile. When people face novel situations they have little if any experience to draw on. In such cases, guessing is not an unreasonable strategy. Similarly, in this model, when a decision maker (regardless of position in the organization) has no experience, it simply guesses, that is, it will randomly respond with a 1 or 0 to each problem. Initially, all situations are novel, thus initially all decision makers are guessing. This results in the organization having initially a 50/50 chance of making a correct decision.

What the members of the organization learn depends on the situation, specifically on what configuration of factors are associated with what decision. For example, if last time an object on the radar was moving quickly and was small it was a missile, the next time that same type of object appears the organization may decide it is a missile whereas if last time a fast small object was a duck the organization learns that fast small objects are not missiles. Similarly, in the proposed model, what the decision makers learn depends on what pattern of 1's and 0's are associated with what type of decision. For the case examined in this paper, the correct answer is a "1" if there are more 1's than 0's in the entire word and a "0" if there are more 0's than 1's. For the problems examined, the word size is odd (27) and there is always a correct decision. Thus, eventually, each analyst learns to be a "majority classifier" for the type of task and organizational structures examined. That is, each analyst learns to simply return a 1 if the majority of the inputs it receives are 1's and 0 if the majority of the inputs it receives are 0's. In the hierarchical structures, upper level managers learn to attend most to those analysts who have a history of producing correct decisions. This adaptive behavior on the part of managers affects which types of tasks are more suited to hierarchies than to teams.

Crisis

A ubiquitous feature of crises is that they are short. Missile display conferences may only be a few minutes (Perrow, 1984). Shuttle launch decisions are made during a few hours or days (Report of the Presidential Commission, 1986). Joint task forces are established for limited engagements (Metcalf, 1986). Should the crisis continue for an extended period, it would not be a crisis but a general problem. The problems associated with the "crisis" would, if they continued, be considered part of the constraints on normal operating conditions. In this paper, this characteristic of crises is captured by making the crisis period short (100 decision periods) relative to the span of decisions made under non-crisis conditions (100,000). The length of the training period relative to the length of the crisis is such that although the individual decision makers continue to learn during the crisis, what they learn has little impact on their performance.

Another important characteristic of crises, as discussed in the introduction, is that communication breakdowns occur and incoming information is often unreliable; for example, computer chips malfunction (Perrow, 1984), maps needed to determine strategy are old (Metcalf, 1986), or people are unavailable (Report of the Presidential Commission, 1986). In order to focus tightly on this aspect of the crisis, only organizations which are engaged in tasks for which they have been trained when the crisis occurs are considered. In this paper, the hypothetical organizations examined, like their real counterparts (such as NORAD and NASA) are during the period of crisis facing the type of problem for which they has been trained but subject to severe operational problems. Clearly one might expect the level of training to affect organizational performance during crisis. Thus, in this paper, only organizations that are fully trained are examined; that is, each organization examined has faced so many problems (100,000) that each decision maker has had experience with multiple examples of each of the sub-problems he or she might face.

Communication breakdowns

During crises, communication breakdowns occur. If a communication breakdown occurs, one or more of the decision makers within the organization become unable to participate in the organization's decision making process for a period of time. When a communication breakdown occurs, decisions cannot flow up the chain of command and feedback cannot flow back down. Firefighters trying to combat a forest fire whose walkie-talkie breaks cannot report on the extent of damage they see, nor can they be told that the fire is spreading to encircle them. Such communication breakdowns differ in at least four ways: (1) why they occur (type), (2) where they occur (location), (3) how long they last (duration), and (4) how many occur at once (severity).

Type: Communication breakdowns occur for a variety of reasons, e.g., radios go silent, people become ill or go on vacation, or are even asleep (Shrivastava, 1987, p. 2). From a structural perspective, what matters to the organization is not the exact reason a breakdown occurs but whether it is a decision maker that is unavailable or a communication channel. That is, the organization may react differently when it is the person that is sick than when it is the telephone that is not working. In Fig. 1 a decision maker being unavailable corresponds to one of the nodes being missing, whereas when a channel is unavailable, one of the lines connecting two nodes in unavailable. As can be seen in Fig. 1, for the centralized hierarchy and the team, whether an analyst (or AEO) becomes unavailable or the channel connecting the analyst to the AEO (or the channel connecting the AEO to the CEO) becomes unavailable, that analyst's (or AEO's) decision will be unknown to the rest of the organization and that analyst (or AEO) will not be able to get feedback. In other words, from an information flow standpoint, for the centralized hierarchy and the team, decision maker and channel unavailability are indistinguishable. In the dual-command hierarchy, on the other hand, since there are multiple channels from and to analysts, having a channel become unavailable is less damaging to the organization than having a decision maker become unavailable.

Location: Communication breakdowns can occur at various levels in the organizational hierarchy. In both the centralized hierarchy and the dual-command hierarchy, decision makers at three different levels can become unavailable – analysts, AEOs and CEOs. In this model, if the CEO in a hierarchy is unavailable, the organization cannot make a decision. Therefore, only cases where either the analysts, the AEOs or both are unavailable are considered. In the team, since there is only one level, only decision makers at that level, i.e. analysts, can become unavailable. Similarly, when the communication channel breaks, the breakdown can occur between analysts and AEOs or between AEOs and CEOs or at both levels. Whether an analyst or the channel between an analyst and an AEO breaks down, I refer to the break as occurring low in the organization. Similarly, whether an AEO or the channel between an AEO and the CEO breaks down, I refer to the break as occurring in the middle of the organization.

Duration: Communication breakdowns vary in length and may not last throughout the entire crisis. For example, prior to the launch of *Challenger*, communication channels were intermittent, i.e. the sites would go on and off mute. Communication breakdowns of three different durations are examined: (1) short – 10 decision periods; (2) moderate – 20 decision periods, and (3) long – 50 decision periods.

Severity: Communication breakdowns also vary in severity. For example, a single phone line can be disconnected or the entire eastern seaboard may be out. Severity can be modeled as the number of breakdowns that occur at the same time. Three different levels of severity are considered: (1) low – one breakdown at a time; (2) medium – two breakdowns at a time; and (3) high – three breakdowns at a time. This range of severity means that between one-ninth and one-third of the organization is unavailable or disabled at any given moment.

Erroneous incoming information

During crises, incoming information may be erroneous or incomplete. For example, defective computer chips cause incorrect information to be displayed, or city fire companies have old maps that don't list current roads and residences; in Bhopal, the police and army helping with the evacuation did not know what type of gas had leaked (Shrivastava, 1987, p. 2). Such errors in incoming information can be modeled by giving the analysts a different set of 1's and 0's than that which occurs in the real problem the organization is facing. For example, if the real problem is "101000000111101101000000111101111," and the first analyst would under non-crisis conditions see the first three bits "101," then during the crisis the first analyst might see the bits "111." Such information errors have a duration and severity just as other forms of communication breakdown. That is, how long the analyst sees incorrect information and how many analysts see incorrect information vary across crises.

Measuring the impact of the crisis

In the previous sections, by identifying characteristics of organizations and crises, idealized "types" have been identified. For organizations, the two parameters – organizational structure (centralized hierarchy, dual-command hierarchy, team) and information access structure (segregated, blocked, distributed) – define nine ways an organization can be designed. Similarly, for crises, six parameters were defined: type (analyst, chanel, incorrect information), location (low, middle, both), duration of the breakdown (short, moderate, long), and severity of the breakdown (low, medium, high). Using simulation, we can analyze how each type of organization fares when faced with each type of crisis. Thus, in effect, the analysis takes the form of a set of "what if" analyses where we repeatedly ask questions of the following forms. If an analyst were unavailable to a team, how much would performance degrade? What if two analysts were unavailable? What if the organization had been structured as a centralized hierarchy instead of a team?

Each parameter is systematically varied so that all possible "what if" questions are asked for each possible organization and each possible crisis. As this is done, the performance of each type of organization given a particular crisis scenario is simulated. In order to mitigate the impact of random fluctuations, a Monte Carlo approach is taken and each type of organization in each type of crisis is simulated 200 times (200 runs). This corresponds to examining 200 different organizations of this type. The results are then averaged to determine the performance of this type of organization when faced with this type of crisis.

Each type of organization is simulated for 100,000 decision periods prior to the onset of the crisis; hence it is faced with a sequence of 100,000 problems (see Fig. 3). Thus, when the crisis occurs, the organization can be viewed as a fully trained organization; that is, each analyst has at that point seen several examples of each class of sub-problem. At this point the organization is then faced with another 100 decision periods of crisis during each of which either a communication breakdown occurs or the incoming information is incorrect. Then, the organization faces another 100 decision periods of non-crisis. This is the recovery period. The specific random sequences for communication breakdowns are repeated across each organizational type in order to contrast how differently designed organizations fare when faced with the same crisis. The specific random sequences for communication breakdowns are varied across the 200 repeated organizations of a particular type so as to prevent bias from a particular random sequence choice.

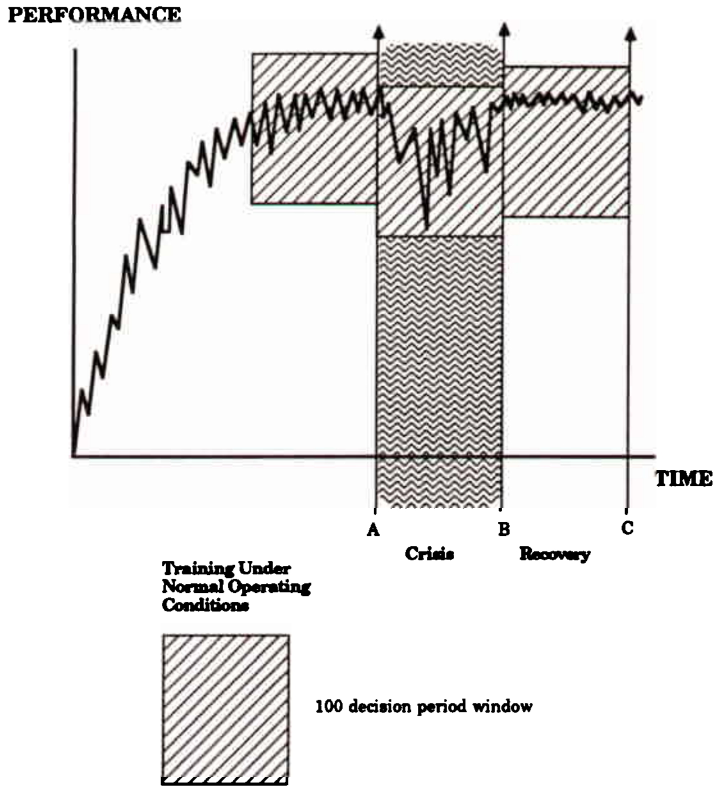

Fig. 3. *Points at which organizational performance is measured.* The dark line indicates the performance of a typical organization over time. The point at which the crisis begins is marked with an A. The point at which the crisis ends and recovery begins is marked with a B. And the point at which recovery is measured is marked with a C. For all of the organizations examined, A equals 100,000; the crisis duration is B – A is 100; and recovery, C – B, is 100 decision periods. It is the ensemble average during the 100 decision period window immediately prior to A, B and C that is used as the measure of pre-crisis, crisis and recovery performance.

To determine the impact of crisis on performance, an ensemble measure of performance in a 100-decision-period window at three different time periods is calculated. These three time periods are: (1) performance under normal operating conditions (100 decision periods prior to onset of crisis); (2) performance under crisis conditions (the 100 decision periods during which crisis is occurring); and (3) performance after crisis (100 decision periods at the end of the crisis). This is illustrated in Fig. 3.

The organization's performance during these 100-decision-period windows is measured as the percentage of correct decisions made by all 200 organizations of that type during those 100 decision periods. A correct decision occurs if the organization's final decision matches the true answer. Since the percentage of correct decisions, denoted by p, is based on the sum of 20,000 binary decisions, the standard deviation of these measures can be determined as:

$$((p(1-6))/20\,000)^{0.5}\times 100$$

for which the maximum value 0.35% occurs when p is 50%.

Performance under normal operating conditions

The impact of a crisis on an organization can only be understood by first examining organizational performance under normal operating conditions. The proposed model predicts that under normal operating conditions, provided there is no turnover in personnel, the organization eventually learns to make the correct decision at some theoretically optimum or "peak" performance level. The fully trained organization operates at this peak performance level. This level is dependent on the type of task the organization is facing and the ratio of incoming information to analysts (Carley, forthcoming). The simpler the task, the greater the consensus, and the less information there is for any member of the organization to process, the more decisions the organization makes correctly. For the task examined in this paper, peak performance is below 100%.

The organization exhibits peak performance when all decision makers, regardless of level, are acting as majority classifiers and so can be thought of as employing as a standard operating procedure, and "propose as their guess about the global majority whatever is in their local majority." Thus, as the organization faces more problems, the members of the organization gain in experience, their performance improves, they stop guessing, and the overall accuracy of the organization's decisions increases. However, the organization does not learn to make 100% of its decisions correctly. There are two reasons for this. First, decision makers pass on decisions, *not* information. Thus, since the decision makers are basing their decision on more than one piece of information, information is being reduced in order to reach a solution. Pertinent information, and general information contrary to the decision, is being lost.[2] Second, the process of channeling and combining information, as happens when information moves from analysts to AEOs to the CEO in a hierarchy, although reducing the complexity of the problem facing each manager, results in yet more information loss. Information loss lowers the peak performance level below 100%.

Such information loss, particularly within hierarchies, is a commonly ob-

[2]Only when the task complexity is nine and the access structure is segregated, and so each analyst sees only one bit of information, is no information lost when the analyst passes on his or her decision. In this paper, however, the task complexity is 27. Thus, regardless of the type of access structure the organizations examined herein suffer information loss.

served phenomenon. March and Simon (1958) point out that as one goes up the hierarchy, uncertainty gets absorbed. For Downs (1967, p. 269), such loss is a function of the condensing of information as it goes up the chain of command. And Jablin (Jablin et al., 1986, pp. 610–613) characterizes the loss in terms of the distortion of information. Such information loss, when it is the result of converting information into decisions, occurs in all organizations regardless of design but to different degrees.

The result is that some types of organizations are more likely than others to make correct decisions even when operating at peak performance (see Fig. 4). First, we see that there is a structural effect; that is, teams generally outperform hierarchies. Teams outperform hierarchies because there is less information loss as decisions are passed up the chain of command. Second, we see that there is a significant information access effect. That is, distributed access is better than blocked access, which in turn is better than segregated access.

Organizations with segregated access structures, where each piece of information is known by only one person, make the most mistakes. This is because analysts know less of the total problem (three pieces of information) and so have less information on which to base their answers.[3] In addition, performance can be improved not only by greater information sharing, as happens in the case of both the blocked and the distributed access structures, but also by

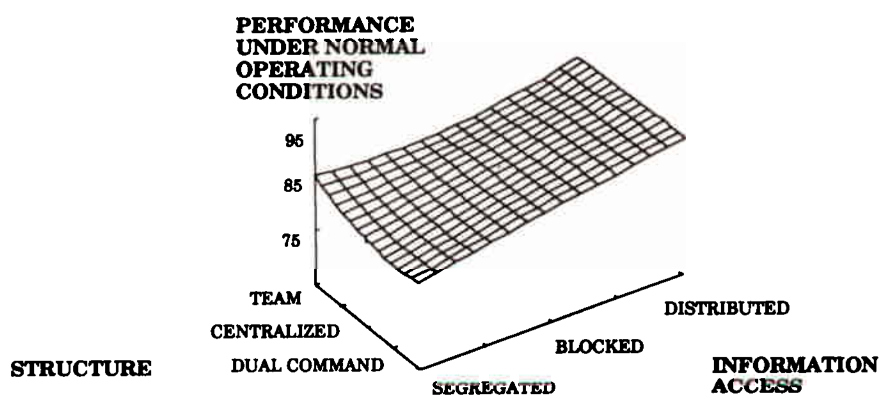

Fig. 4. *Under normal operating conditions, distributed teams perform the best.* This plot shows the performance level for organizations under normal operating conditions as a function of the organizational structure and the information access structure. This surface was generated using a 3-D negative exponential interpolation.

[3]In another study (Carley, 1990), it was found that information sharing did not always improve performance, particularly if one was concerned with the rate of learning or performance after so many time steps rather than performance in a fully trained organization as was examined in this paper.

carefully arranging who knows what information. Blocking information, such that a group of decision makers sees identical information (while it clearly promotes consensus among that group, a consensus that allows them to form a coalition), is less beneficial to the organization than distributing that same information across a wider group. The distributed access structure promotes the highest performance because each analyst knows more and the organization as a whole gains better resolution on a problem if each piece of information is examined from multiple perspectives (that is, in conjunction with different sets of other information).

There is also an interaction between access and structure. When information access is distributed or segregated, teams outperform centralized hierarchies, which in turn outperform dual-command hierarchies. However, when a blocked access structure is used, the performance of all structures is almost identical. That teams generally outperform hierarchies is accounted for by the greater information loss as one goes up the chain of command in the hierarchy as previously described. When a blocked structure is used, however, this effect is mitigated as three different analysts all have identical information and so present the same decision. Even though there are nine analysts, there are only three different decisions. In the centralized hierarchy, all of the analysts working for one manager agree, so the manager's decision is the same as the analysts'. Thus, the team and centralized hierarchy behave identically. In the dual-command hierarchy, most of the analysts working for one manager agree, so the manager's decision tends to be the same as the analysts'. Thus, in general, the dual-command hierarchy makes the same decision as the team or the hierarchy, although now and then the dual-command hierarchy will make more correct decisions, even when information is blocked.

When crises occur

Crises, regardless of the type of organization or the specific features of the crisis, tend to degrade performance (see Fig. 5). But after the crisis, the organization returns to its pre-crisis performance level or sometimes does even a little better. The correlation between pre- and post-crisis behavior is 0.996. Post-crisis performance improvements occur as the organizations are adaptive and their personnel continue to learn even during crises. Further, since even during the crisis the organization is doing the task for which it is trained, crisis lessons continue to be relevant and will be useful in the future, and during future crises. An example of this type of learning was noted by Perrow (1984, p. 287) when he noted that three days after the false alarm incident described at the beginning of this paper, "... the identical alarm recurred. SAC crews again started

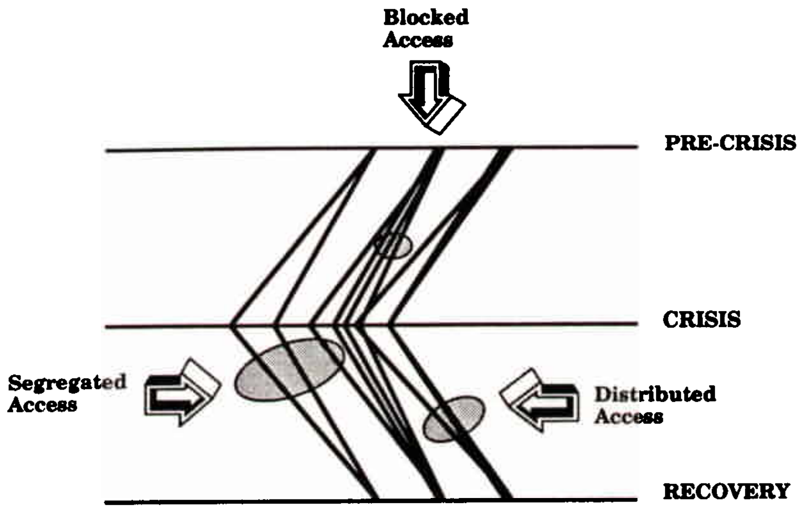

Fig. 5. *Crises temporarily degrade organizational performance.* For all organizations, performance decreases during crisis and then during the recovery period following the crisis returns to pre-crisis levels. In this figure, performance level increases as one moves from left to right. Each of the lines crossing the three periods represents the average behavior for all organizations with that coordination structure. Within each of the information access structures, the organizations are (going from left to right) the centralized hierarchy, the dual-command hierarchy and the team.

their engines, and again the alarm was determined to be false in less than three minutes. NORAD switched to a backup computer, and began the search for a possible malfunction." For NORAD, lessons learned during the first crisis affected later behavior.

Although organizational performance degrades due to the crisis, not all types of organizations are equally affected. In particular, centralized hierarchies are much more affected by crises than either dual-command hierarchies or teams (see Fig. 6). And organizations with a distributed access structure are more affected than organizations with a segregated access structure, which are more affected than organizations with a blocked access structure (see Fig. 7). Due to this variation in resiliency to crisis, teams and dual-command hierarchies perform the best during the crisis, whereas centralized hierarchies with segregated information access perform the worst during crises. What this is suggesting is that during crises, organizations such as NORAD which have more of a dual-command distributed structure are likely to outperform organizations such as NASA/Thiokol which use more of a dual-command segregated structure and which in turn are expected to outperform organizations such as Union Carbide which use more of a centralized segregated structure.

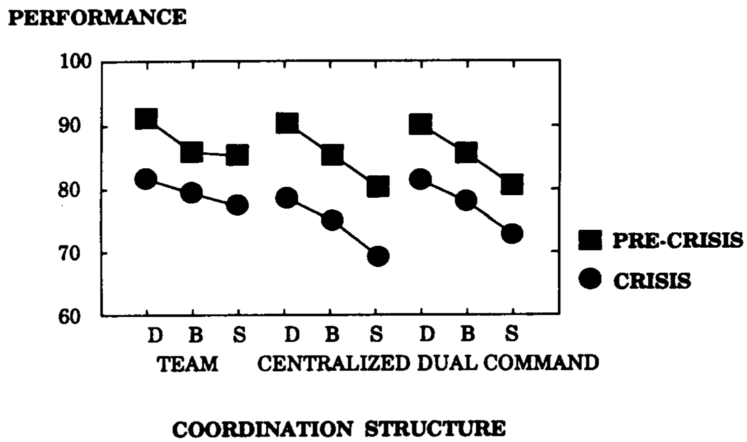

Fig. 6. *Teams generally outperform hierarchies.* The pre-crisis (square) and crisis (circle) performance level for all organizations with a particular coordination structure are shown. Going from left to right, within each organizational structure, as the information access structure varies from distributed (D) to blocked (B) to segregated (S), the performance of the organization, whether under normal operating conditions (pre-crisis) or crisis conditions (crisis), drops. Each dot is the average performance value for all organizations in that category.

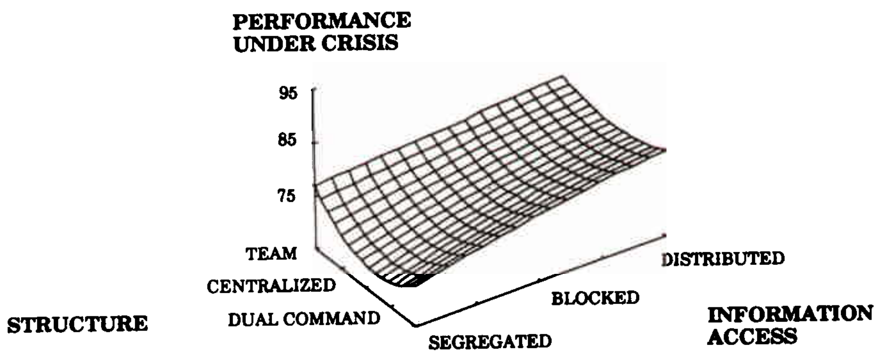

Fig. 7. *Under crisis conditions, distributed access in the best.* This plot shows the performance level for organizations under crisis conditions as a function of the organizational structure and the information access structure. This surface was generated using a 3-D negative exponential interpolation.

What type of crisis is most debilitating?

Of the types of crises examined, those in which the incoming information is erroneous generally degrade organizational performance more than those in which either decision makers become unavailable or communication channels break down (see Fig. 8). Only when the information access structure is segregated is it worse to have channel or decision makers unavailable than to have

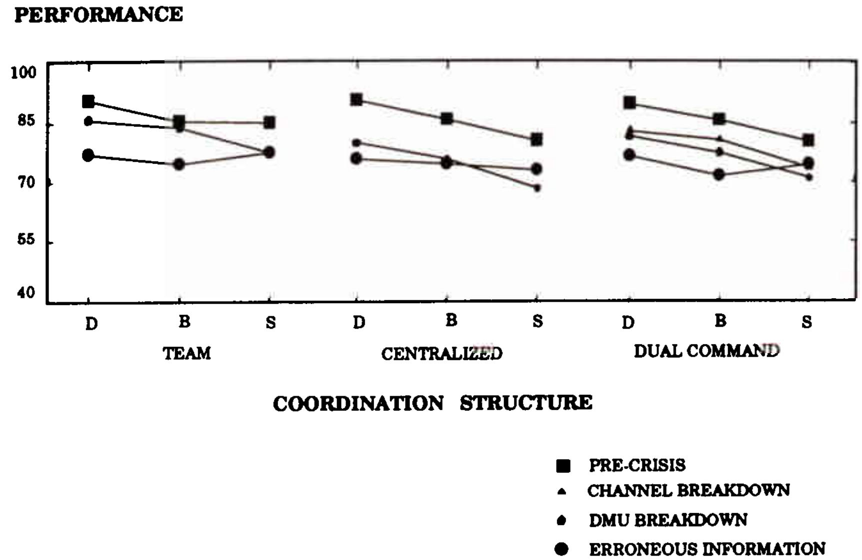

Fig. 8. *Erroneous information is generally more debilitating than breakdowns.* Each dot is the average performance value for all organizations in that category for the indicated type of crisis. For teams and centralized hierarchies, crises where the decision maker becomes unavailable and those where the channel breaks cannot be distinguished.

erroneous incoming information. This result follows from the fact that when there is redundant access to incoming information, if the information is erroneous there are more decision makers who are misled by it. But this same redundancy enables the organization to maintain higher levels of performance even when decision makers become unavailable and communication channels break down.

For teams and centralized hierarchies, channels breaking down and decision makers becoming unavailable are functionally identical events. Only in the dual-command hierarchy do the two events produce different results, and in this case channel breakdowns are the least debilitating. For the dual-command hierarchies, performance degrades more when analysts become unavailable than when channels break down, as even if a channel breaks there are alternate routes by which the information can move up the hierarchy. In contrast, when the decision maker is unavailable information is lost.

Duration

For the various crises examined, different decision makers are incapacitated, either through a breakdown or through errors in incoming information, for various amounts of time. For the organizations examined, the length of a par-

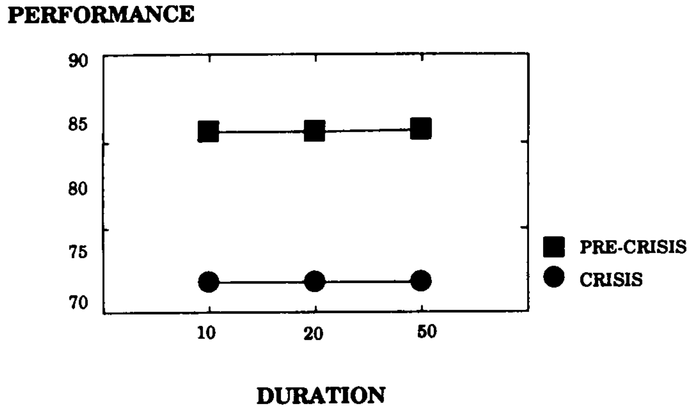

Fig. 9. *Duration of breakdowns does not affect performance.* For a fully trained organization, the length of time that a particular decision maker or channel is unavailable, or the length of time that the information coming into a particular analyst is wrong, does not affect organizational performance. Each dot is the average performance value for all organizations in that category across all types of crisis.

ticular incapacitation has little impact on performance (see Fig. 9). Note that the issue is not how long the crisis lasts, as all crises examined are the same length. During a crisis, one or more decision makers are always incapacitated. What we desire to know is whether it is worse for the organization for one decision maker to be incapacitated for the entire crisis period or for a set of decision makers, one after the next, to be incapacitated. The answer is that it doesn't matter. When the organization is fully trained there is little local adaptation during the crisis. Consequently, it doesn't matter which decision maker is unavailable, or whose phone is out, or whose information is incorrect. It does matter, however, how many decision makers are simultaneously incapacitated, and the level of those decision makers in the organizational hierarchy also matters. For example, at NORAD it didn't matter whether it was the SAC display monitor, the satellites or the radar that indicated incoming missiles; what mattered was the number of cites that indicated incoming missiles. Since it was only one cite, they knew the information was erroneous and so could avoid making a wrong decision. For the shuttle, the decision process would have been the same whether it was Thiokol or KSC that was offline, but when both went offline the decision process changed drastically.

Severity

As the severity of the crisis increases, the number of decision makers who are

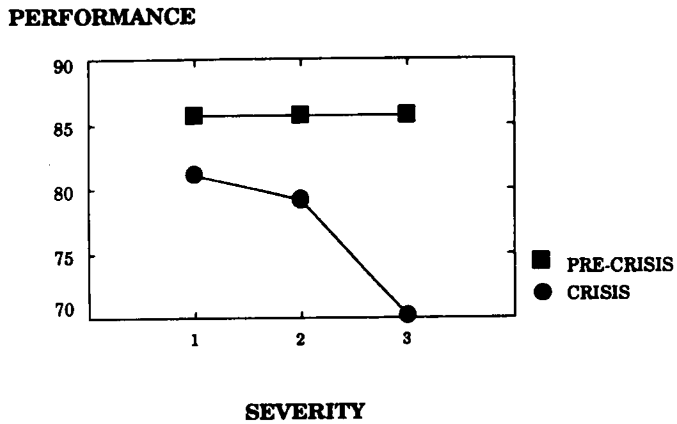

Fig. 10. *Increasing severity decreases performance.* The severity fo the crisis increases as the number of decision makers whose performance is simultaneously compromised increases (whether due to unavailability, channel breakdown or erroneous incoming information). Going from left to right, the severity index (1, 2, 3) indicates the number of "sights" that simultaneously break down. For decision maker breakdowns, this is the number of decision makers that are not communicating. For channel breakdowns, this is the number of channels that cannot be used. For erroneous incoming information, this is the number of analysts who are receiving incorrect new information. Each dot is the average performance value for all organizations in that category across all types of crisis.

incapacitated at the same time increases. Not surprisingly, the more severe the crisis, the more the organization's performance is decreased from its pre-crisis level (see Fig. 10). This is true regardless of the organizational coordination scheme or the type of crisis. Or, in other words, the fewer decision makers available, and the less accurate information the organization has to go on, the more likely the organization is to make mistakes. At NORAD, if both the radar and the SAC display monitor indicated incoming missiles, the chance of a counterattack being called for would be more likely.

Location

During the crisis, any of the decision makers within the organization may be compromised. Both analysts and managers can take ill, and phone lines or other communication channels at any level can break. The proposed model suggests that when the problems occur in the middle of the organization the results are more devastating than when they occur at the lowest level (see Fig. 11). Specifically, in hierarchies, performance is least degraded by the crisis when only analysts are compromised, and most degraded when AEOs are compromised. This is true regardless of the type of crisis or the type of hierarchy or the information access structure. Further, when comparing teams with the hierarchies when only analysts are compromised, teams still outperform hierarchies.

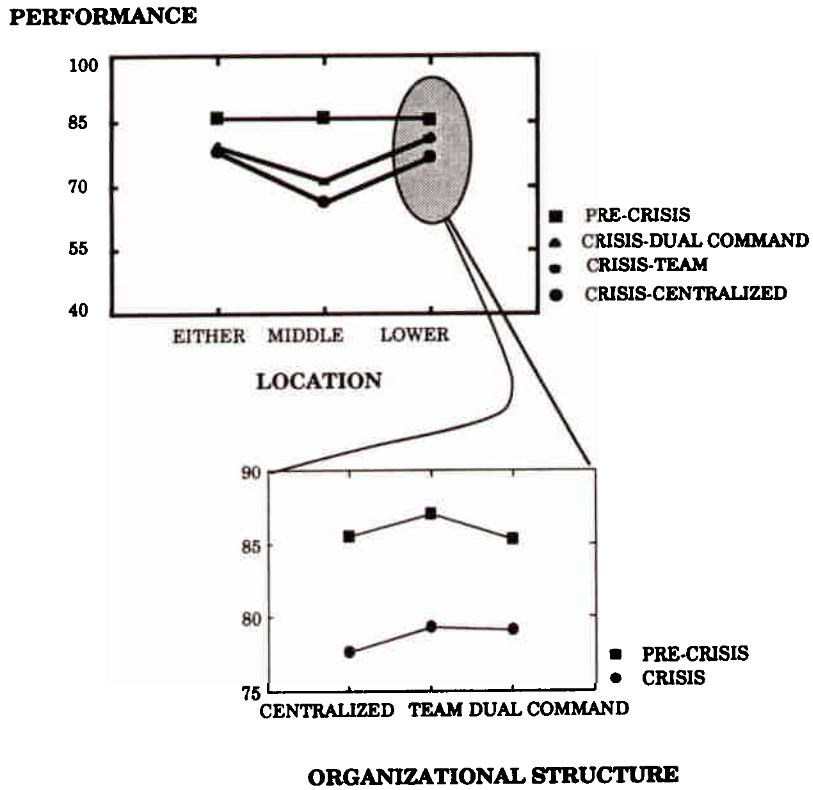

Fig. 11. *Breakdowns higher in the organization have more drastic consequences.* Displayed in the top graph is the performance of hierarchies as the location of the crisis in the organization is varied, as both analysts and AEOs are compromised (left), just AEOs (middle), and just analysts (right). In the bottom graph, pre-crisis (squares) and crisis (circle) level performance for teams and hierarchies given that only analysts are compromised is shown. Each dot is the average performance value for all organizations in that category.

In Fig. 11 it can be seen that if both analysts and managers are compromised, the organization actually fares better than if just managers are compromised. This result is due to the way in which breakdowns were modeled. Regardless of where in the organization the error occurred, the severity and duration of crisis were held constant. Thus imagine that one decision maker was unavailable for 10 periods. If this person is at the low level then an analyst is unavailable. If this person is at the middle level then an AEO is unavailable. If this person is at "either level" then the missing person is sometimes an analyst and sometimes an AEO.

Discussion

Crises, in this paper, have been characterized as periods of limited duration during which there are one or more communication breakdowns or the value of incoming information is suspect. While communication breakdowns and erroneous information are clearly aspects of crisis, there are other aspects as well. Other aspects of crises include a rapid increase in the rate of incoming information, a rapid increase in the rate at which decisions must be made, a sense of uniqueness, and the potential creation of a new command structure. And, of course, all of these features could and often do occur at the same time and may have synergistic effects.

The first two of these aspects, increase in the rate of incoming information and increase in the required decisions, indicate the importance of timing in analyzing organizational behavior during crises. In the model proposed herein, there is no absolute sense of time; rather, time is divided into a series of decision periods. Were measures of absolute time incorporated, issues such as delay due to communicating decisions and feedback through different channels could be examined. Such an examination, might provide a very different view of which organizational structure is the most effective during crisis. For example, Carley et al. (1988) found that distributed structures, as opposed to hierarchies, could respond more quickly to crises but were more likely to make mistakes.

For many organizations, facing a crisis means facing a situation for which the organization is not prepared or for which there has been little actual experience. Some crises are, after all, quite rare. In such cases, experience gained under normal operating conditions may not be transferable, thus resulting in an organization whose personnel have no experience. Similarly, when a crisis occurs, some organizations create completely new organizational structures to deal with that crisis. For example, when a natural disaster occurs, the Red Cross, in conjunction with local service groups, constructs a special temporary organization to cope with that disaster. In such new temporary organizational structures, the individuals have never worked together before, let alone on this particular type of crisis. Whether due to novelty of the task (crisis) or the organizational structure, the group of individuals who must make decisions and integrate their behavior in an attempt to cope with the crisis come into the situation with either no experience, or different levels of experience, or inappropriate experience, not to mention differences in norms and procedures for behaving. In this paper, the members of the organization all began the crisis with not only the same level and type of experience, but they were fully trained. And it was from such a highly prepared state that the ability of the organization to cope with crisis was determined. As a consequence, what was learned during the crisis had little effect on performance and so post-crisis performance was almost identical to pre-crisis performance, and the length of the crisis and the duration of the communication breakdowns and other incapacitations were irrelevant to behavior. Although not examined herein, the proposed model could be used to determine which type of organizational structure can best cope with crises when the personnel are less optimally prepared. Such an analysis might lead to different conclusions than those drawn herein. In particular, such an examination would develop our understanding or learning and adaptive learning during the crisis and so would allow an examination of questions such as, "Do crises promote learning the wrong thing?"

Additionally, individual and organizational behavior have also been oversimplified. Many of these simplifications have been discussed elsewhere (Carley, forthcoming) and need not be repeated here. Given the focus on crisis, attention is now centered on those simplifications of individuals and organizations that are particularly important when studying crisis.

Individuals within the organization are modeled as perfect historians engaged in experiential-based decision making. Admittedly, this is an oversimplification of actual human decision making behavior; these "individuals," despite being intendedly adaptive and imperfect statisticians, do not exhibit all of the cognitive limitations known to affect individual decision making behavior, particularly under stress. Yet, under stress, individuals tend to make more mistakes, take longer to solve problems, are less able to solve problems, and satisfice. However, in the proposed model, individuals exhibit the same problem solving and information processing skills during crisis as they do under normal operating conditions. Thus, future work should consider the additional impact of such compromises to individual cognition that might result during crises. A second aspect here is that individuals as perfect historians rely as much on information that they received a year ago as that which they receive today. Yet there are saliency effects such that individuals attend more to the most recent information. Were the decision makers to attend more to recent information, the value of training might be decreased, with the upshot that the organization is even more affected by crisis than suggested in this study. Factors such as stress and salience, should they be taken into account, might effect different relative performance among the types of organizations examined.

Organizations are treated as having relatively fixed structures; that is, the number of decision makers is fixed, there is no turnover, and even if decision makers become incapacitated these same analysts later re-emerge and take over their former positions, and so on. Yet during crises, personnel may continue to turn over. For example, during a crisis the organization may not know that a communication breakdown is temporary and may promote someone to fill the gap or, in an effort to "right the wrong," personnel may be fired or transferred.

Such turnover, unlike the temporary communication breakdowns that were examined in this paper, creates an organization in which the personnel may come into the crisis with different levels and types of experience. Further, since turnover tends to make organizations learn slower and less (Carley, forthcoming), organizations who must cope with turnover may be less affected by crisis than those examined herein where there was no turnover.

Despite its limitations, the proposed model does capture many of the features of crisis and organizational behavior. These include information loss, communication breakdowns, information sharing, limited rationality on the part of decision makers, and organizational structure. With some modification, the proposed model can serve as a framework to look at many of the issues just discussed. As an example, we might consider implementing turnover with different ratios of executive-to-staff turnover, or even different types of turnover such as a performance-based or a tenure-based model of turnover in order to examine whether turnover augments the effect of communication breakdowns on organizational performance during crisis. As a final example, this framework could be used to look at the impact of the timeliness and accuracy of feedback on organizational learning and hence performance. In the studies presented in this paper, organizations received prompt and accurate feedback. While organizations learn from experience, the feedback that is provided may be incomplete, slow, inaccurate and subject to interpretation, particularly during periods of crisis. Thus, future studies should consider the role of incomplete or slow feedback.

Conclusion

Policy implications

A fairly general model of adaptive organizational performance has been proposed and its implications for organizational performance under normal operating conditions and crisis examined. A variety of policy implications follow from this analysis. Several of these will be considered. First, however, a word of caution. These policy implications are only for fully trained organizations and for organizations that routinely face the same type of problem. If an organization is not fully trained or if every problem faced by an organization is unique, then different policies might be suggested. While these implications may seem obvious in retrospect, it is important to note that very few organizations follow all of these policies.

The first implication is that when multiple project teams, or multiple departments, must work together to solve a particular problem and the organization has the resources so that multiple personnel have access to the same information, then personnel who share the same information should be in different teams or departments rather than blocked within a single team or department. Admittedly, having all members of group know the same information increases consensus within the group and reduces the reliance of the manager of any one group member. This makes the group more able to withstand personnel turnover and communication breakdowns. However, what is good for the group is not good for the organization. When information is shared within the group, but not between groups, each group effectively forms a coalition with a unified perspective on the problem that can be completely different from the perspective of other groups. Consequently, the likelihood of agreement across groups is reduced and the likelihood of incorrect decisions is increased. When information is blocked within a group, in the aftermath of a crisis, comments such as, "if we had known that ... we would have decided differently," should frequently occur, comments such as this would also be common when no information is shared by personnel within the organization (the segregated access structure). For example, prior to the *Challenger* accident, NASA, Marshall and Thiokol were basically using a blocked information access structure. In the aftermath of the accident, this blockage was pointed to as a key determinant of the accident. As reported by Aldrich (shuttle program director), reviews conducted the previous summer on the Solid Rocket Booster project between NASA and Marshall were not brought through his office. Were such information brought to his office, concern over O-ring erosion may have been noticed. In contrast, when information is distributed across groups, such a scenario is less likely to result. At NORAD, the information access structure is more distributed than blocked. Consequently, information on the same air space is available on multiple channels – such as satellite and radar. Thus, false alarms, as in the example at the beginning of the paper, are relatively easy to locate. A further advantage of distributed structure is that it prevents catastrophic information loss. That is, during a crisis, should a communication breakdown occur that makes an entire group unavailable, more information is lost when information is blocked within a group than when it is distributed across groups. In addition to these consequences, which are predicted by the proposed model, such blocked structures may result in the formation of coalitions that engender interorganizational strife and may increase the competition and quarreling between groups.

The second implication is that organizations should expend more effort in acquiring correct information than in setting up extra communication channels. In other words, having the latest in communication equipment and making it possible for, and encouraging, all personnel to communicate to all other personnel regardless of group boundaries will not help the organization make better decisions if the information the organization has to act upon is incorrect. For example, local fire companies often have a very limited budget and old equipment. When a fire occurs, they need to get to the building that is on fire, yet city maps are often out of date. Given that they have a limited amount of money to spend, should they spend it on keeping their own records as to where buildings are located or should they spend it on walkie-talkies for communicating during the fire? This study suggests that in the absence of good city maps, the local fire company should keep its own records.

The third implication is that organizations should spend more effort making sure managers are available and their communication channels intact than making sure staff and other workers (and their communication channels) are available. The reason is that, within hierarchies, the closer the communication breakdown is to the top of the pyramid the more drastic the consequences. This suggests higher pay for executives to prevent turnover, making sure the president's phone lines work before the staff's, and so on.

Although every organization must face a crisis at some point or another, the following question remains. Should organizations structure themselves so that when a crisis does occur it will have little effect? The answer to this question is particularly important if, as suggested by Starbuck et al. (1978), when faced by crises, organizations re-emphasize their existing structure rather than switch to a new form or, as suggested by Zucker (1984, p. 5), alternative designs become literally unthinkable. The answer, however, depends on the frequency with which the organization expects to face crises and the cost of these crises when they do occur. For all fully trained organizations, having some information redundancy is a good thing; that is, information redundancy improves performance both under normal operating conditions and during crisis. However, only a few organizations should consider employing redundant communication channels.

Let us consider the cost of the different coordination schemes. Coordination schemes in which multiple personnel have access to the same information, such as in the blocked and the distributed structures, improve performance in a fully trained organization. Such redundant information access schemes are costly from an information storage and processing perspective as each analyst has more to do. In contrast, redundant communication channels, such as those that occur in the dual-command structure, rarely improve the performance of the fully trained organization. Redundant communication channels are costly from both an information processing and a communication standpoint. Further, the information processing costs are higher when communication channels are redundant than when access to incoming information is redundant, as the additional processing occurs higher up in the administrative hierarchy. Given that redundant communication channels have great additional costs and may not improve performance, the only organizations that may find them efficacious are those that expect frequent crises or for which the cost of a crisis, however unlikely, is extreme. For other organizations, the increase in communication and information processing costs may not be recoverable.

But what of the less fully trained organization or the organization which is facing a problem it has never faced before? Redundancy, when it increases the complexity of a task faced by an individual decision maker, slows the rate at which the organization learns (Carley, 1990). In an untrained organization, given multiple personnel access to the same information or having the same information being sent up to multiple managers via multiple communication channels will impair performance. Should a crisis occur before the organization is fully trained, the organization may actually be worse off than it would have been had it used a simpler coordination scheme. An organizational design that works well for a highly trained unit like NORAD may actually be a liability for an organization that cannot afford to operate at such a highly trained level.

References

Anderson, P.A. and Fischer, G.W., 1986. A Monte Carlo model of a garbage can decision process. In: J. March and R. Weissinger-Baylon, eds. Ambiguity and Command: Organizational Perspectives on Military Decision Making. Pitman, Boston, MA: 140–164.

Arrow, K.J. and Radner, R., 1979. Allocation of resources in large teams. Econometrica, 47: 361–385.

Bar-Shalom, Y. and Tse, E., 1973. Tracking in a cluttered environment with probabilistic data association. In: Proceedings of the Fourth Symposium on Nonlinear Estimation, University of California, San Diego, CA: 13–22.

Bavelas, A., 1950. Communication pattern in task-oriented groups. Journal of the Acoustical Society of America, 22: 730–735.

Bond, A. and Gasser, L., 1988. Readings in Distributed Artificial Intelligence. Kaufmann, San Mateo, CA.

Burns, T. and Stalker, G., 1961. The Management of Innovation. Tavistock, London.

Carley, K.M., 1986. Efficiency in a garbage can: implications for crisis management. In: J. March and R. Weissinger-Baylon, eds. Ambiguity and Command: Organizational Perspectives on Military Decision Making. Pitman, Boston, MA: 196–231.

Carley, K.M., 1990. Coordinating for success: trading information redundancy for task simplicity. In: Proceedings of the 23rd Annual Hawaii International Conference on System Sciences, Vol. 3: 261–270.

Carley, K.M., forthcoming. Organizational learning and personnel turnover. In: Organization Science, to appear.

Carley, K.M., Lehoczky, J., Rajkumar, R., Sha, L., Tokuda, H. and Wang, L., 1988. Comparing approaches for achieving near optimal solutions in a distributed decision making environment. Working paper. Carnegie Mellon University, Pittsburgh, PA.

Cohen, A.M., 1962. Changing small-group communication networks. Administrative Science Quaterly, 6: 443–462.

Cohen, A.M., Robinson, E.L. and Edwards, J.L., 1969. Experiments in organizational embeddedness. Administrative Science Quarterly, 14: 208–221.
Cohen, M.D., March, J.G. and Olsen, J.P., 1972. A garbage can model of organizational choice. Administrative Science Quarterly, 17(1): 1–25.
Crecine, J.P., 1986. Defense resource allocation: garbage can analysis of C3 procurement. In: J.G. March and R. Weissinger-Baylon, eds. Ambiguity and Command: Organizational Perspectives on Military Decision Making. Pitman, Marshfield, MA: 72–119.
Cyert, R.M. and March, J.G., 1963. A Behavioral Theory of the Firm. Prentice-Hall, Englewood Cliffs, NJ.
Davis, S. and Lawrence, P., 1977. Matrix. Addison-Wesley, Reading, MA.
DeGroot, M.H., 1974. Reaching a consensus. Journal of the American Statistical Association, 69: 118–121.
Downs, A., 1967. Inside Bureaucracy. Little, Brown & Co., Boston, MA.
Galbraith, J., 1973. Designing Complex Organizations. Addison-Wesley, Reading, MA.
Gloves, T. and Ledyard, J., 1977. Optimal allocations of public goods: a solution to the free-rider problem. Econometrica, 45: 738–809.
Hart, G. and Goldwater, B., 1980. Recent False Alerts from the Nation's Missile Attack Warning System: Report to the Senate Committee on Armed Services. US Government Printing Office, Washington, DC.
Hastie, R., 1986. Experimental evidence on group accuracy. In: 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: 129–158.
Hastie, R., Park, B. and Weber, R., 1984. Social memory. In: R.S. Wyer and T.K. Srull, eds. Handbook of Social Cognition. Erlbaum, Hillsdale, NJ.
Hinnings, C. and Greenwood, R., 1988. The Dynamics of Strategic Change. Blackwell, New York.
Jablin, F.M., Putnam, L.L., Roberts, K.H. and Porter, L.W., eds., 1986. Handbook of Organizational Communication: an Interdisciplinary Perspective. Sage, Beverly Hills, CA.
Johnson, M.K. and Hasher, L., 1987. Human learning and memory. Annual Review of Psychology, 38: 631–668.
Kemeny, J.G., 1981. Report of the President's Commission on the Accident at Three Mile Island. Pergamon Press, New York.
Kimberly, J., 1987. The study of organization: toward a biographical perspective. In: J. Lorsch, ed. Handbook of Organizational Behavior. Prentice-Hall, Englewood Cliffs, NJ.
LaPorte, T.R. and Consolini, P.M., 1988. Theoretical and operational challenges of high reliability organizations: air traffic control and aircraft carriers. Working paper. University of California, Berkeley, CA.
Levitt, B. and March, J.G., 1988. Organizational learning. Annual Review of Sociology, 14: 319–340.
Lichtenstein, S. and Fischhoff, B., 1977. Do those who know more also know more about how much they know? The calibration of probability judgments. Organizational Behavior and Human Performance, 20: 159–183.
Lindblom, C.E., 1959. The "science" of muddling through. Public Administrative Review, 19: 79–88.
Malone, T.W., 1986. A formal model of organizational structure and its use in predicting effects of information technology. Technical report. Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA.
Malone, T.W., 1987. Modeling coordination in organizations and markets. Management Science, 33: 1317–1332.

March, J.G. and Olsen, J.P., 1975. The uncertainty of the past: organizational learning under ambiguity. European Journal of Political Research, 3: 147–171.

March, J.G. and Simon, H.A., 1958. Organizations. Wiley, New York.

Marschak, J., 1955. Elements for a theory of teams. Management Science, 1: 127–137.

Masuch, M. and LaPotin, P., 1989. Beyond garbage cans: an AI model of organizational choice. Administrative Science Quarterly, 34: 38–67.

Metcalf, J., 1986. Decision making and the Grenada rescue operation. In: J.G. March and R. Weissinger-Baylon, eds. Ambiguity and Command: Organizational Perspectives on Military Decision Making. Pitman, Marshfield, MA: 277–297.

Michael, E.J., 1986. Elements of effective contingency planning. In: Avoiding and Managing Environmental Damage from Major Industrial Accidents, Proceedings of an International Conference. Air Pollution Control Association, Pittsburgh, PA.

Padgett, J.F., 1980a. Managing garbage can hierarchies. Adminstrative Science Quarterly, 25(4): 583–604.

Padgett, J.F., 1980b. Bounded rationality in budgetary research. American Political Science Review, 74: 354–372.

Panning, W.H., 1986. Information pooling and group decisions in nonexperimental settings. In: 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: 159–166.

Perrow, C., 1972. Complex Organizations: a Critical Analysis. Scott, Foresman & Co., Glenview, IL.

Perrow, C., 1984. Normal Accidents: Living with High Risk Technologies. Basic Books, New York.

Radner, R., 1987. Decentralization and incentives. In: T. Groves, R. Radner and S. Reiter, eds. Information, Incentives and Economic Mechanisms: Essays in Honor of Leonid Hurwicz. University of Minnesota Press, Minneapolis, MN: 3–47.

Report of the Presidential Commission on the Space Shuttle Challenger Accident, 1986 (William, P. Rogers, Chair). US Government Printing Office, Washington, DC.

Rumelhart, D., McClelland, J. and the PDP Research Group, 1986. Parallel Distributed Processing: Explorations in the Microstructure of Cognition. MIT Press, Cambridge, MA.

Shaw, M.E., 1981. Group Dynamics: the Psychology of Small Group Behavior. McGraw-Hill, New York.

Shrivastava, P., 1987. Bhopal: Anatomy of a Crisis. Ballinger, Cambridge, MA.

Sills, D.L., Wolf, C.P. and Shenaski, V.B., eds., 1982. Accident at Three Mile Island: the Human Dimension. Westview Press, Boulder, CO.

Simon, H.A., 1947. Administrative Behavior. Macmillan, New York.

Smith, R., 1980. The contract net protocol: high-level communication and control in a distributed problem solver. IEEE Transactions on Computers, C-29(12): 1104–1113.

Starbuck, W., Greve, A. and Hedberg, B., 1978. Responding to crises. Journal of Business Administration, 9: 111–137.

Steeb, R., Cammarata, S., Hayes-Roth, F. and Wesson, R., 1980. Distributed intelligence for air fleet control. Technical report. Rand Corporation.

Steinbruner, J.D., 1974. The Cybernetic Theory of Decision Process. Princeton University Press, Princeton, NJ.

Strand, R.G., 1971. An efficient suboptimal decision procedure for associating sensor data with stored tracks in real-time surveillance systems. In: Proceedings of the IEEE Conference on Decision and Control, Miami Beach, FL: 33–37.

Thorndyke, P., McArthur, D. and Cammarata, S., 1981. Autopilot, a distributed planner for air fleet control. In: Proceedings of the Seventh IJCAI, Vancouver, BC: 171–177.

Tsitsiklis, J.N. and Athans, M., 1984. On the complexity of decentralized decision making and detection problems. IEEE Transactions on Automatic Control, AC-30: 440–446.

Tversky, A. and Kahneman, D., 1971. The belief in the law of small numbers. Psychological Bulletin, 76: 105–110.

Tversky, A. and Kahneman, D., 1974. Judgment under uncertainty: heuristic and biases. Science, 185: 1124–1131.

Weber, M., 1922. Bureaucracy. In: H. Gerth and C.W. Mills, eds. Max Weber: Essays in Sociology. Oxford University Press, Oxford: 196–244.

Wilensky, H., 1967. Organizational Intelligence: Knowledge and Policy in Government and Industry. Basic Books, New York.

Zucker, L., 1984. Organizations as institutions. In: S.B. Bacharach, ed. Research in the Sociology of Organizations, Vol. 2. JAI Press, Greenwich, CT: 1–47.