Coordinating for Success: Trading Information Redundancy for Task Simplicity

Kathleen Carley
Associate Professor of Sociology
Department of Social and Decision Sciences
Carnegie Mellon University
1990

In Proceedings of the 23rd Annual Hawaii International Conference on Systems Sciences

Kona, Hawaii

Coordinating for Success: Trading Information Redundancy for Task Simplicity

Kathleen M. C. Carley

Department of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

This paper examines whether the coordination scheme used in the organization, i.e., the level of redundancy in access to information and the organizational structure which defines who commands/communicates-to whom, affects the organization's ability to learn and hence its ultimate performance. A model of organizational decision making is presented where organizational performance is dependent on the combined decisions of the individual decision makers who base their decisions on their previous experience. Using simulation the impact of organizational structure, information redundancy, and personnel turnover on organizational performance as the organization faces a sequence of similar, but not identical problems, is explored. This research suggests that increasing redundancy in information access does not necessarily compensate for personnel turnover, and may actually decrease the rate of organizational learning and degrade performance. The explanation lies in the fact that increasing redundancy may also increase the complexity of the task faced by individual decision makers. Thus, the adage *divide and conquer* may actually define the coordination scheme that for the organization produces the greatest success.*

1. Introduction

Larry's Diner is a small but growing enterprise. At first, Larry could "do it all" - take the order, fry the hamburgers and fries, ... ring up receipts. As business expanded, Larry needed to hire help, eventually to the point of hiring another cook. Initially both cooks cooked entire orders. They found that they were constantly getting in each other's way. So Larry divided up the cooking job and hired additional cooks so that one could do the frying, another the pies, another the salads, and so on. In streamlining, or moving to the "assembly line method", the complexity of the task faced by each cook was reduced. The difficulty, however, was that should one cook be ill the others did not necessarily know what to do thus resulting in events like the pie kitchen shutting down.

The tradeoff, between information redundancy and task complexity, that this story illustrates pervades the core of organizational activities. Different structural solutions exist in different organizations ranging from the critical employee, such as the archetypical secretary who is so completely relied on by her boss that should she quit the office would fall apart, to the complete subdivision of tasks that theoretically characterizes military and unionized activities. From another vantage point we see that redundancy and task complexity correspond to different types of costs that the organization must face. For example, redundancy corresponds to "storage" costs and task complexity to "processing" costs. From this vantage point it becomes clear that the redundancy/task-complexity tradeoff, and related concerns, also arise in the design of information systems, decision support systems, and expert systems. In all cases, institutional memory (or data) is distributed such that what the organization knows is a function not only of which individual decision maker knows what but the relationships (such as communication and command channels) between these decision makers. Thus, coordinating the organization in order to ensure a particular level of task performance entails both finding the right balance between redundancy (in terms of multiple agents having access to, and needing to process, the same information) and the level of task complexity faced by a single agent (in terms of the amount of information available and the resultant complexity of the information processing task) and placing the right links or communication channels between personnel. In human, as well as computer organizations, finding the right balance of redundancy and complexity is further complicated by the fact that people leave, systems go offline, etc.

This paper addresses two related questions in the area of organizational coordination: (1) how does the choice of a particular level of redundancy and task complexity affect the organization's performance given that the organization is composed of intelligent but boundedly rational decision makers, and (2) are certain coordination schemes better than others given that personnel do turnover. The coordination scheme is defined as the C3I structure (Command, Communication, Control, and Intelligence), associated order of processing, and rules for processing which in this paper is characterized by the organizational structure (who commands/communicates-to whom, the procedure for producing a final organizational decision), and the access structure (who has access to what information and hence the level of information redundancy). These questions are addressed by examining via simulation the performance

*This research was supported by the NSF under grant No. SES-8707005.

over time of organizations of intelligent agents which differ in the scheme they use to coordinate these agents.

Broadly speaking, within the area of organizational coordination, three different paradigms have emerged. The first paradigm, organized anarchies (also referred to as "garbage can theory") [1; 2; 3; 4], views organizational decisions as emerging, often through oversight, as a result of a matching up of energy, choices, and problems as personnel, choice situations, and information flow through the organization in a highly volatile fashion. This perspective emerged from a tradition in which the organization's behavior is seen as affected by the intendedly, but boundedly, rational behavior of the individual decision makers within the organization [5; 6; 7; 8; 2]. Under this paradigm, in the extreme case, there is no coordination. The second paradigm, distributed decision making [9; 10; 11; 12; 13; 14; 15; 16]views organizational decisions as being constructed by personnel working together in a cooperative fashion on tasks where the nature of the task limits information access and the quality and type of information available. This paradigm has emerged primarily from the artificial intelligence tradition in which the individual's behavior is seen as affected by the available information, search mechanisms, and plans. Under this paradigm, coordination is critical as the tasks are too complex for single individual decision makers to handle and coordination emerges from the dynamic exchange of information between peers. In contrast to these paradigms, structural theorists from the classical [17] to the neo-classical [18] argue that the structure of the organization (both formal and informal [19; 20]) controls organizational behavior. Under this paradigm, coordination is not problematic and there exists a "best" structure independent of the limitations on individual behavior. This paper draws from these diverse paradigms to suggest that decisions are constructed in a coordinated environment in which the flow of personnel, choices, and information constrain the effectiveness of the coordination scheme. In addition, it is argued that the extant organizational structure as a coordination structure affects the performance of the individual decision makers and consequently the behavior of the organization not only because the structure limits who has access to what information but also because different structures are differentially effective at storing and retrieving information. This paper, like much previous research, explores organizational behavior as a function of individual behavior; but, unlike previous research various organizational structures are contrasted (specifically the centralized hierarchy and the distributed team).

2. Experiential Learning Model

The organization is engaged in a quasi-repetitive integrated decision making task.** Thus, there is a sequence of decision making periods such that each period the organization faces a new problem which is similar, but not identical, to previous problems. During each period, the new problem is evaluated by the decision making units (DMUs), a decision is made, and the members of the organization are informed of the "correctness" of their decision for that problem. As the organization, and each individual decision maker, sees a sequence of problems it builds up a typical response pattern - i.e. it learns.

2.1. Classification Task

The particular task looked at, determining whether there are more 1's than 0's in a binary word of length N, is a classification task. That is, the organization given a particular problem must classify it as a problem that has "more 1's" or as a problem that has "more 0's". The complexity of the task environment (i.e. problem space) is defined as the length of the words (N). The length of a word is simply the number of bits or positions in that word that can be 1 or 0. When there are N positions there are 2^N distinct words (i.e. problems). The complexity of the task environment does not change over time. In this paper, the level of task complexity examined is $N = 27$.***

Each decision period the organization is faced with a particular problem which is divisible into a set of subproblems. A problem is a word of length 2^N. A subproblem is a portion of the word. Each analyst is given a single subproblem. All analysts see the same size subproblem. The size of subproblems and the degree of overlap in who knows what depends on the level of redundancy as defined by the information access structure.

In order to examine the impact of the task on performance two versions of this task will be examined. In version one, the unbiased task, the problems will be drawn from the full set of 2^N words with replacement. All problems are equally likely and the probability of a 0 being the correct decision (50%) is equal to the probability of a 1 being the correct decision (50%). In version two, the biased task, the distribution of problems will be slightly skewed so that the probability of a 0 being the correct decision is slightly greater than the probability of a 1 being the correct decision (54%/46%).

2.2. Coordination Scheme

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

**A task is quasi-repetitive if 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 the other by non-repetitive tasks. A task would be repetitive if exactly the same problem is faced over and over again; whereas, a task is non-repetitive if the problem is unique. 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 making units (DMUs) within the organization.

***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 9 analysts. When there is no redundancy and 9 analysts the possible choices of task complexity are odd multiples of 9 (9, 27, 45, 63, ...). When redundancy is admitted the task complexity must be greater than 9. Thus, 27 is the simplest case that meets these criteria.

2.2.1. Two Organizational Structures

Two organizational structures are examined: the centralized hierarchy and the distributed team (see figure 1). The primary difference between these two structures is the presence of upper level management in the hierarchy and the absence of such management in the team. The presence of such upper level management, by mediating the decisions made by lower level DMUs (analysts), may potentially reduce the impact of turnover thus affecting the level of redundancy needed for equivalent performance. These structures are not meant to exhaust the set of potential or actual organizational structures. Rather, they represent two idealized structural types which are interesting due to their prevalence in real organizations and which because of their structural difference (presence of upper level management) we expect to be differentially affected by turnover. Both the hierarchy [17; 21; 3; 22]and the team [1; 23; 24; 25; 26; 27; 28; 29]have been extensively studied, but their performance has rarely been contrasted. In addition, these two structures represent different ends of the organizational spectrum in terms of the degree to which institutional memory is centralized.

Centralized Hierarchy: The centralized hierarchy is modeled as a three tier organization composed of a chief executive officer (CEO), a set of assistant executive officers (AEOs), and a set of analysts. Each analyst in each decision period receives information (a subproblem), 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, makes the final integrated decision (yes or no), finds out if it is correct, and then informs each AEO of the correct final decision. Then each AEO informs each analyst of the correct final decision. It is from these decisions and the resultant feedback that the analyst's experience is formed. In this paper, the specific centralized hierarchy examined has 13 DMUs with 3 under each "manager" as in figure 1. There are 9 analysts. A hierarchy of 13 DMUs is the minimum size non-trivial hierarchy that can be examined such that the hierarchy has 3 levels and an odd number of DMUs under each "manager".

Distributed Team: The distributed team is modeled as a single tier organization composed of a set of analysts. Each analyst, each decision period, receives information (a subproblem), 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. The analysts then find out the correct decision. In this paper, the specific team analyzed has 9 DMUs all of which are analysts as in figure 1. The team has 9 DMUs in order to match the number of analysts in the hierarchy. The number of analysts, rather than total DMUs, is matched so that the complexity of the subproblem seen by analysts, regardless of structure, is identical given a task complexity and redundancy level.

Both Structures: The analyst, given his or her subproblem, must decide - yes ("I think there are more 1's than 0's in the full problem" represented by a 1) or no ("I think there are more 0's than 1's in the full problem" represented by a 0). Thus, each decision maker is making a recommendation for what he or she thinks the final decision should be. The individual decision maker by passing on a 1/0 decision rather than the number of 1's has compressed information; hence, there is information loss. Information loss is higher the more complex the subproblem faced by the DMU and the greater the number of organizational levels.

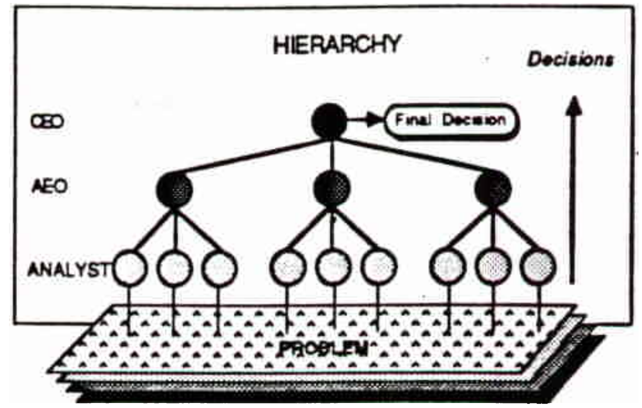

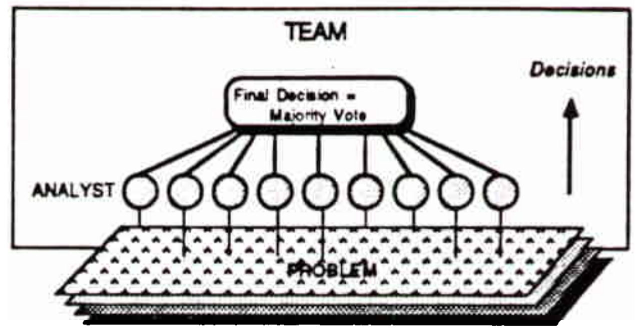

Figure 1: Organizational Structures

Two organizational structures are examined - the centralized hierarchy (top) and 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 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. Thus, in the centralized hierarchy, institutional memory is centralized in the upper management, albeit in a reduced information form. In the distributed team, the analyst's decision is his "vote" as to what the final decision should be. In the distributed team the majority vote is the final decision. In the distributed team institutional memory is completely distributed. In both cases it is only the analysts who have access to the "raw data" associated with the problem.

Problem solution requires integrating the analysts' decisions. For each problem there is a decision provided by the organization (the final decision) and a true decision. The final decision provided by the organization is, for the hierarchy, the decision made by the CEO and is a "1" if the CEO decides there are more 1's than 0's in the problem, and a "0" if the CEO decides there are more 0's than 1's in the problem. For the distributed team, the final decision is the majority vote made by the analysts - "1" if more analysts think there are more 1's or "0" if more analysts think there are more 0's. For both hierarchies and teams the correct answer is a "1" if there really are more 1's than 0's in the problem and "0" if there really are more 0's than 1's. For the problems examined, the word size is odd (27) and there is a true decision. The organization's decision is correct if the final decision matches the true decision.

2.2.2. Information Access Structure

Who knows which pieces of information, and hence the level of redundancy, is defined by the information access structure. The level of information redundancy is defined as the average number of analysts who know each piece of information. To illustrate the division of the problem into subproblems given different levels of redundancy let us imagine an organization with 3 analysts which is faced with the problem - 101010001. In the no-redundancy case (where each piece of information is known by only one analyst) each analyst has a distinct subproblem which is a contiguous set of positions and each position in the word is evaluated by only one analyst. For the illustrative organization and problem the first 3 positions (101) would be the subproblem assigned to one analyst, the second three positions (010) would be assigned to a second analyst, and the last three positions (001) would be the subproblem assigned to the third analyst. Now consider the same organization with a low level of redundancy (1.67). For the illustrative organization and problem the first 5 positions (1-5) (10101....) would be the subproblem assigned to one analyst, the five positions 4-8 (...01000.) would be assigned to a second analyst, and the five positions 1-2 and 7-9 (10....001) would be the subproblem assigned to the third analyst.

Redundancy is implemented by treating the problem, which is a bit string, as a circle and giving each analyst a contiguous block of bits such that their first bit occurs in the position where it would occur if there was no redundancy. Under this implementation: a redundancy level of 1 means that each analyst has access to 3 pieces of information, each piece of information is known by only one analyst, and there is no overlap in who knows what; a redundancy level of 1.67 means that each analyst has access to 5 pieces of information, there is 1 piece of information known by each analyst that is not known by some other analyst; a redundancy level of 2.33 means that each analyst has access to 7 pieces of information, each piece of information is known by at least two analysts, and no analyst has access to information that no one else has; and so on. When the redundancy level equals the number of analysts (9), we have a situation where all analysts have access to all information which corresponds to the case of complete information/complete overlap discussed in earlier studies [22; 1]. As the level of redundancy increases the complexity of the subproblems faced by the analysts increase. In this paper, four levels of redundancy will be examined - 1, 1.67, 2.33, and 3. These levels are chosen so that the size of the subproblem seen by the analysts is always odd - 3,5,7,9.

2.3. The Decision Makers

All DMUs, regardless of type (analyst, AEO, or CEO), are experiential based decision makers. Each DMU keeps a cumulative record of the subproblems that it receives, its decisions, and the true answer. For each DMU each subproblem that it sees falls into a particular class. A class is a particular pattern of 1's and 0's, such as 010. For example, in the no-redundancy case since each analyst sees three positions an analyst will have 2^3 or 8 classes of problems and 0.125 of all subproblems seen by that analyst will be in each class. As the DMU encounters subproblems it builds up, for each class of subproblems, 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 DMU's experience that, given this class of problems, the true decision was a 0. The expectation that the answer is a 1 is defined as the proportion of times in this DMU's experience that, given this class of problems, the true decision was a 1. When the DMU is faced with a subproblem the first thing the DMU does is to *determine what class the problem is in*. This is a simple pattern matching process and simply takes longer the more complex the subproblem (the more bits). Then, once the DMU knows what class of problems it is working on it makes a decision using the following heuristics:

1. If the expectation of a 0 is greater than the expectation of a 1 return 0 as the decision.
2. If the expectation of a 0 is less than the expectation of a 1 return 1 as the decision.
3. 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.

Within the organizations examined, each DMU (regardless of position in the organization) will initially have no experience to draw on and so will randomly respond with a 1 or 0 to each problem. The organization will thus initially have a 50/50 chance of making a correct decision. Eventually, each analyst, faced with either the unbiased or biased task will learn to be a "majority classifier". That is, each analyst will learn to simply return a 1 if the majority of the inputs it receives are 1's and a 0 if the majority of the inputs it receives are 0's. This is true for both the hierarchy and the team. The upper level management in the hierarchy, however, will learn different behaviors depending on the task. Thus, differences in task performance for hierarchies and teams are due solely to what the upper level management learns.

2.4. Personnel Turnover

Organizational turnover occurs when members of the organization leave and new personnel need to be hired. When turnover occurs the organization loses the expertise of the DMU who leaves and gains the expertise of the DMU who joins the organization. Among the characteristics of turnover that determine its impact on the organization's ability to learn are the rate of turnover and the level and type of experience possessed by the new employees.

Turnover is viewed as a continual process where one analyst leaves, another enters, another leaves, another enters throughout time. Turnover is implemented by having an analyst leave the organization, and another immediately enter the organization periodically over time as a Poisson process. Which analyst leaves the organization is determined randomly; all analysts are equally likely to be chosen to leave. The rate of turnover is defined as one over the mean number of decision periods between these exits/entrances (mean inter-arrival time). Five turnover rates are examined: (1) no turnover - 0.0; (2) low - 0.01 - every 100 decision periods , (3) medium-low - 0.02 - every 50 decision periods, (4) medium-high - 0.05 - every 20 decision periods, and (5) high - 0.10 - every 10 decision periods. In this paper, all newly hired analysts are novices; i.e., they have no information on any subproblem.

2.5. Analysis Procedure

In order to examine how the performance of an organization changes overtime given continuous turnover and continuous learning on the part of the individual decision makers Monte-Carlo simulation is used. In the foregoing discussion four parameters were identified: structure, turnover rate, redundancy, and task type. By varying the value of these parameters 40 types of organizations for each task type can be specified. Each type of organization is simulated 400 times. This corresponds to examining 400 different organizations of this type. Each organization is simulated for 2500 decision periods (hence it is faced with a sequence of 2500 problems). The random sequences for both turnover and problem choice are not repeated across runs nor across organizational types in order to prevent bias from a particular random sequence choice. Although the hierarchy and team each see a different set of 400 sequences of 2500 problems this does not affect the results as these sets are drawn from the same underlying distribution. The same argument is true with respect to turnover.

For a particular type of organization, performance at a specific time is measured as the percentage of correct decisions made by all 400 organizations of that type. A correct decision occurs if the organization's final decision matches the true answer. Two specific measures of learning will be used - the "final level of learning" or "final performance" and the rate of learning. The final performance is defined as the percentage of correct decisions made in the last 200 decision periods (periods 2300 to 2500) by all 400 organizations. The percentage of correct decisions is an estimate of the ensemble probability of a correct decision at that decision period. For most organizations examined performance has plateaued prior to the 2000th decision period. Thus, by averaging the last 200 periods together a better estimate of final performance is achieved. For final performance the standard deviation is $((p(1-p))/80000)^{.5} \times 100$ which will always be less than .0.18%. The rate of learning is defined as the average number of decision periods it takes until the organization has increased its performance by 10% (learned to make 60% rather than 50% of its decisions correctly). For those organizations that never learn to make 60% of the decisions correctly the rate of learning is defined to be 1250 (half the maximum number of time periods). Typical standard deviations are in the range 10 to 20.

3. Can Redundancy Mitigate Turnover?

In the organizations examined, generally, as the rate of personnel turnover increases the performance decreases. The impact of information redundancy is less clear sometimes mitigating and sometimes enhancing the impact of turnover. Regardless of the rate of turnover or the level of redundancy. both teams and hierarchies, generally learn to make more correct decisions.

3.1. Structure and Performance

Let us begin by looking at organizational learning as a function of organizational structure (no turnover, no information redundancy). The expectation is that, over time, the members of these organizations will learn and their performance will improve. This expectation is born out. Both hierarchies and teams learn; however, which type of organization learns more depends on the task.

With an unbiased task teams rapidly come to outperform hierarchies. Teams, in only 35 (σ = 4.59) decision periods, learn to make 60% of the decisions correctly; whereas, hierarchies take 195 (σ = 15.74) decision periods to learn to make 60% of their decisions correctly. The final level of performance for teams is 85% (σ = 0.12%) and for hierarchies it is 80% (σ = 0.14%). Teams outperform hierarchies simply because, in making the final decision, teams have access to more information (all analysts get a vote). In contrast, in hierarchies there is more information loss as the CEO does not see all of the analysts' decisions but a set of reduced decision made by the AEOs which in turn are based on the reduced information provided by the analysts. Since many of the problems are not decomposable (i.e., there are problems where the distribution of 1's and 0's are such that the correct decision can not be recovered from the subproblems) neither teams nor hierarchies learn to make 100% of their decisions correctly. Hierarchies, however, suffer greater information loss than teams and so learn less. In general, teams learn faster and better than hierarchies when the task is unbiased. As the level of redundancy changes, as the level of turnover increases, teams typically learn more than hierarchies. Across all hierarchy team pairs, where for each pair the level of turnover and redundancy is the same, the average difference in the mean performance for teams and hierarchies is 2.6% (σ = 0.11%).

With a biased task hierarchies come to outperform teams. The final level of performance for teams is 67.5% (σ = 0.17%) and for hierarchies it is 70.7% (σ = 0.16%). Like the unbiased task, however, hierarchies still learn slower than teams: the learning rate is 35 (σ = 3.73) for hierarchies as opposed to 15 (σ = 1.73) for teams. Hierarchies come to outperform teams when the task is slightly biased as the team has no upper level management to learn about the bias. Recall that the analysts in both teams and hierarchies learn the same thing, to be majority classifiers. The overall decision for the team is made by majority rule and so there is no mechanism for responding to bias. Thus, teams, despite less information loss fare worse than hierarchies as the individual decision makers do not have enough information to learn that the task is biased. In contrast, in hierarchies, despite the information loss the upper level management can effectively piece together more information and hence come to recognize that the task is biased. As will be seen in the next sections, hierarchies do not always outperform teams when faced with biased tasks.

3.2. Personnel Turnover and Performance

Now consider the impact of personnel turnover where all new employees are novices, there is no information redundancy, and the task is unbiased. In both hierarchies and teams, the higher the rate of turnover the lower the final performance level (see figure 2). This is a monotonic relation. Performance is directly and negatively affected by turnover. In hierarchies, as the rate of turnover increases from 0 to 0.1 final performance drops from 80.5% (σ = 0.14%) to 63.1% (σ = 0.17%). Similarly, in teams, as the rate of personnel turnover increases from 0 to 0.1 final performance drops from 85.0% (σ = 0.13%) to 64.7% (σ = 0.17%). In both cases performance drops off rapidly at first as the rate of turnover increases. Personnel turnover degrades organizational performance because it is a perpetual drain on resources and portions of the institution's memory leaves as personnel leave. The perpetual training of new personnel necessitated by turnover detracts from the organization's ability to improve its performance.

In contrasting hierarchies and teams we see that for the same level of turnover teams learn faster and better, and hence outperform, hierarchies. The difference in final performance between teams and hierarchies decreases as the rate of turnover increases. For example, there is a difference of 3.3% (σ = 0.21%) between teams and hierarchies when the rate of turnover is low (0.01) and there is a difference of 1.6% (σ = .24%) between teams and hierarchies when turnover is high. This suggests that the higher the

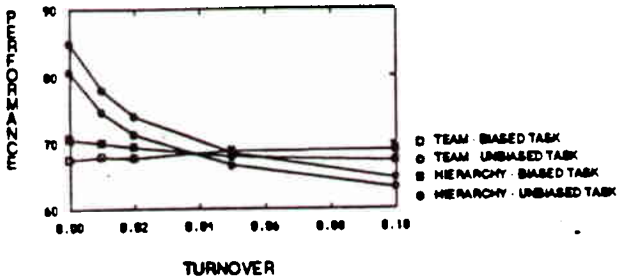

Figure 2: Turnover Degrades Performance When the Task is Unbiased, But Has Little Effect When the Task is Biased

Final performance decreases for both teams (white circles) and hierarchies (black circles) as the rate of turnover increases when the task is unbiased. When the task is biased (squares) teams (white squares) actually show a slight improvement in performance as turnover increases. All new personnel are novices (no prior experience) and there is no information redundancy. *Personnel turnover*, going from left to right, increases from 0.0 (no turnover) to 0.10 (approximately one person leaves and another enters every 10 decision periods). Each dot represents the averaged behavior of 400 organizations of that type over 200 time periods.

expected level of turnover the less important the choice of organizational structure, but for low levels of turnover the team is a preferable structure.

In contrast, when the task is biased, turnover has little effect on final performance (see figure 2). In hierarchies, as the rate of turnover increases from 0 to 0.1 final performance drops from 70.7% (σ = 0.16%) to 67.3% (σ = 0.17%). Whereas, in teams, as the rate of personnel turnover increases final performance slightly increases from 67.5% (σ = 0.17%) to 68.9% (σ = 0.16%). As we saw in the last section, when the task is biased analysts do not have enough information to determine that the task is biased and so will act as they would when the task is unbiased. Thus, personnel turnover has less effect when the task is biased as the analysts are less effective to begin with.

In addition, as the rate of personnel turnover increases the rate at which organizations learn decreases for hierarchies but not for teams (see figure 3). In hierarchies, as the rate of turnover increases from 0 to 0.1 the rate of learning decreases which is seen in that the number of decision periods it takes to learn to make 60% of the decisions correctly increases from 195 (σ = 11.17) to 595 (σ = 257.19). Whereas, in teams, as the rate of turnover increases the rate of learning remains at 35. In contrast, when the task is biased hierarchies can actually learn faster (see figure 3). In hierarchies, the rate of learning is not monotonically related to the rate of turnover and instead changes from 35 (σ = 3.73), to 25 (σ = 2.90), to 35 (σ = 4.05), and back. Whereas, for teams faced with a biased task, as the rate of turnover increases from 0 to 0.1 the rate of learning remains at 15. Contrasting figures 2 with 3 we see that organizations learn faster when the task is biased but they learn less, regardless of their structure unless the rate of turnover is quite high.

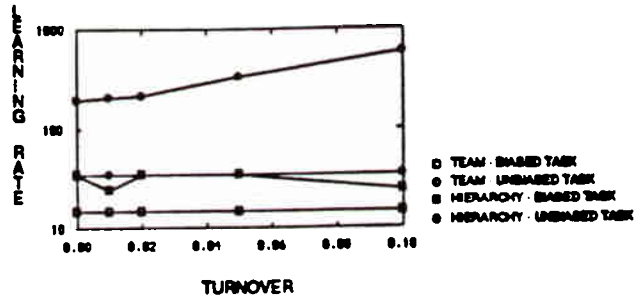

Figure 3: The Rate at Which Hierarchies, but Not Teams, Learn is Affected by Turnover

Hierarchies learn slower the higher the turnover when faced with an unbiased task (black circles), but may learn faster when faced with a biased task (black square). The rate at which teams learn (white dots) , on the other hand, is not affected by turnover. The number of decision periods until organizations learn to make 60% of their decisions correctly for hierarchies (top line) and teams (bottom line) is displayed. The higher this value the slower the organization learns All new personnel are novices (no prior experience) and there is *no information redundancy*. *Personnel turnover*, going from left to right, increases from 0.0 (no turnover) to 0.10 (approximately one person leaves and another enters every 10 decision periods). Each dot represents the averaged behavior of 400 organizations of that type.

3.3. Information Redundancy and Performance

Now let us consider the impact of information redundancy on organizational performance. First, consider organizations where there is no turnover and the task is unbiased. In both hierarchies and teams, information redundancy degrades final performance from performance under no redundancy conditions (see figure 4). For example, in hierarchies the final performance level changes from 80.5% (σ = 0.14%) when there is no redundancy, to 80.0% (σ = 0.14%) at the 1.67 redundancy level, to finally end at 76.9% (σ = 0.15%) when the redundancy level is 3. Similarly, in teams the final performance level changes from 85.0% (σ = 0.13%) when there is no redundancy, to 84.3% (σ = 0.13%) at the 1.67 redundancy level, to finally end at 79.4% (σ = 0.14%) when the redundancy level is 3. In dramatic contrast, when the task is biased and there is no turnover redundancy can actually improve performance, but it is not guaranteed to do so (see figure 4). For example, for hierarchies final performance increases from 70.7% (σ = 0.16%) when there is no redundancy to 73.9% (σ = 0.16%) when the redundancy level is 3. And for teams, final performance increases from 67.5% (σ = 0.17%) when there is no redundancy to 77.4% (σ = 0.15%) when the redundancy level is 3. When the task is biased both teams and hierarchies perform better with some information redundancy. In the hierarchy, high levels of redundancy actually begins to degrade performance, whereas for teams performance keeps improving as redundancy increases.

In addition, redundancy has a mixed effect on the rate at which organizations learn. With just a little redundancy the organizations learn faster; but, as the level of redundancy increases the rate of learning actually slows down (see figure 5). When the task is biased, the higher the redundancy the slower the teams learn; however, for hierarchies redundancy has little if any effect on the rate of learning (see figure 5).

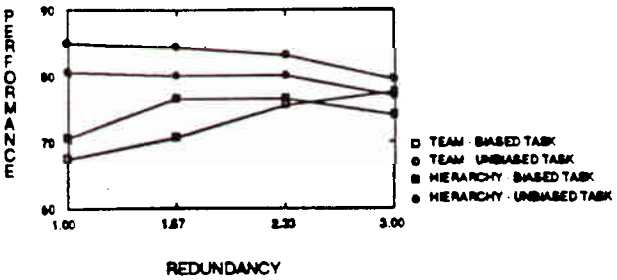

Figure 4: Redundancy Has a Mixed Effect on Performance

When the task is unbiased (circles) the performance of both teams (white dots) and hierarchies (black dots) degrades as redundancy increases. But, when the task is biased (squares) as redundancy increases the performance of teams (white dots) improves and the performance of hierarchies (black dots) first improves and then degrades. Redundancy, going from left to right, increases from 1 (no redundancy) to 3 (each piece of information is seen by 3 people). Each dot represents the averaged behavior of 400 organizations of that type for 200 time periods.

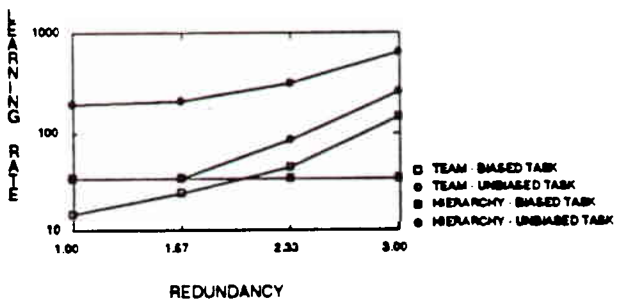

Figure 5: Organizations Learn Slower the Greater the Redundancy

Except for the case of hierarchies faced with a biased task (black squares) the rate at which the organization learns decreases as the level of redundancy increases. The number of decision periods until organizations learn to make 60% of their decisions correctly for hierarchies (black dots) and teams (white dots) is displayed. The higher this value the slower the organization learns. Redundancy, going from left to right, increases from 1 (no redundancy) to 3 (each piece of information is seen by 3 people). Each dot represents the averaged behavior of 400 organizations of that type.

These "mixed effects" due to redundancy follow from the fact that when information redundancy is increased competing forces are brought to bear on the problem. These competing forces are complexity (each analyst by having access to more information must deal with more classes of subproblems and so learns slower) and resolution (each analyst by having access to more information can make a better estimate of the true answer). In the limit, in the absence of turnover, when there is complete redundancy (each analyst sees the entire problem) these competing forces would lead to learning being extremely slow, but perfect. When the task is biased, redundancy tends to improve performance as the greater complexity means that the analysts have more information and so are more able to learn that the task is biased. For teams, which in the absence of redundancy fail to learn that the task is biased, redundancy actually improves performance but at the cost of slowing the rate of learning. For hierarchies, which through upper level management were already capable of learning that the task was biased the greater complexity due to redundancy has less of a positive effect.

3.4. Combining Redundancy and Turnover

Now let us consider the combined effect of redundancy and turnover When the task is unbiased the main effect, for both hierarchies and teams, is that rather than mitigating the effect of turnover, redundancy and turnover reinforce each other to the point that when both are high the organizations' performance is only slightly better than chance. The greater task complexity, resulting from greater redundancy, means that decision makers are learning less which when coupled with more learned personnel leaving, when turnover is higher, results in increasingly worse performance. A more minor effect is that, regardless of organizational structure, when the rate of turnover is high a low level of redundancy can actually mitigate the impact of turnover and improve performance. When turnover is high, performance is so low that the greater resolution provided by redundant information comes in to play prior to the complexity effect. These effects are illustrated in figure 6 where the final performance of hierarchies is graphed as a function of both turnover and redundancy. Although not shown, teams exhibit identical behavior with the exception that they have a higher level of performance.

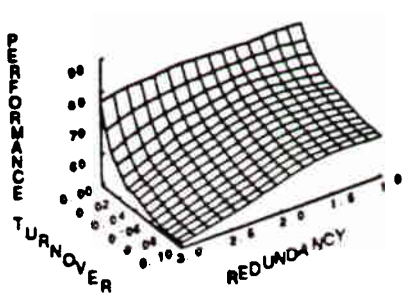

Figure 6: For Hierarchies When the Task is Unbiased Redundancy and Turnover work Together to Degrade Performance

This plot shows the final level of learning (and hence performance) as a function of the rate of personnel turnover and the level of redundancy for hierarchies. All new personnel are novices (no prior experience). Low levels of redundancy somewhat mitigate the decrease in performance due to turnover - this is seen in the upward bend in the graph on the front right. This mitigation varies by the level of turnover as can be seen in the fact that as the level of turnover increases the positive benefit of redundancy decreases. This surface was generated using a 3-D negative exponential interpolation procedure [30; 31].

In contrast, when the task is biased, hierarchies and teams differ in their behavior. For hierarchies low levels of redundancy can mitigate low levels of turnover (see figure 7) but with high levels of turnover redundancy is irrelevant. Low levels of redundancy improve performance due to increased resolution. Contrasting figures 6 and 7 we see that when the task is biased the hierarchies performance is less affected overall by turnover and redundancy than it is when the task is unbiased.

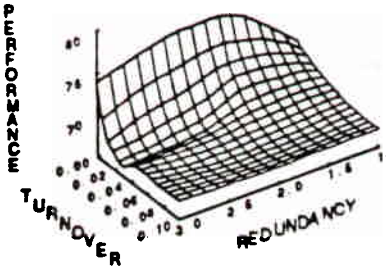

Figure 7: For Hierarchies When the Task is Biased Redundancy can Mitigate Turnover

This plot shows final performance for hierarchies as a function of the rate of personnel turnover and the level of redundancy when the task is biased. All new personnel are novices (no prior experience). Low levels of redundancy somewhat mitigate the decrease in performance due to turnover - this is seen in the upward bend in the graph on the front right. This mitigation varies by the level of turnover as can be seen in the fact that as the level of redundancy increases there is a dip in performance when turnover is low, but there is no such dip when turnover is high. This surface was generated using a 3-D negative exponential interpolation procedure [30; 31].

For teams, when the task is biased. redundancy can mitigate or enhance the effect of turnover such that there is an optimum level of redundancy per turnover level (as seen by the ledge before the fall in figure 8). When redundancy and turnover are both high, rather than mitigating the effect of turnover, redundancy and turnover reinforce each other to the point that the organizations' performance is only slightly better than chance - final performance is 56.6%. As with hierarchies, when the task is biased overall performance is less affected by turnover and redundancy than it is when the task is unbiased.

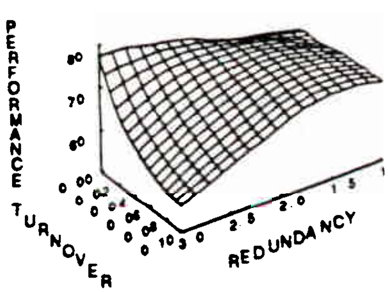

Figure 8: For Teams When the Task is Biased There is an Optimum Redundancy Level per Turnover Level

This plot shows the final performance level for teams as a function of the rate of personnel turnover and the level of redundancy. The ledge, at an angle from low turnover to low redundancy, shows that there is an optimum level of redundancy per rate of turnover. All new personnel are novices (no prior experience). This surface was generated using a 3-D negative exponential interpolation procedure [30; 31].

4. Discussion

The model proposed is an over simplification of individual and organizational behavior. In particular, the DMUs are very limited agents and the organization is facing a very simple task environment. Individual decision makers 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 nor are they capable of more complex forms of learning such as rule creation. In the proposed model there are no framing effects as the individual always correctly categorizes a problem which may reduce the number of incorrect decisions made by the individual decision makers hence producing greater organizational performance. And, in the proposed model there are no saliency effects as the individual correctly remembers all experiences and makes decision using all previous experiences. Yet, the performance of human decision makers is affected both by how they frame problems [32] and what information they consider salient (or are able to recall) [33; 34]. Similarly, problem framing and information saliency or availability affects organizational performance [35; 36]. Further, the inability to create new rules - such as responding with a 1 as the final answer with a certain probability if its expectation is greater than the expectation of a 0 - results in an inability to determine that the task is biased on the part of the decision makers. In addition the individual decision makers are cooperating to locate the true decision. By focusing on cooperative behavior the effect of such forces as differential goal setting, negotiation, bargaining, persuasion, and gaming are missing from this formulation. And, it is assuredly the case that within organizations such behaviors occur [37]. Finally, since the organization is facing only one task at a time there is no question of who should attend to what problem nor do the decision makers become confused by having multiple goals or having too much information as they might were they trying to solve multiple problems at once. Collectively these various factors probably result in the organizations examined learning faster and more, and hence coming to perform better, than would organizations composed of more limited agents in a more complex less cooperative task environment.

Now let us consider other features of the model such as size, turnover, and redundancy. In this study, all organizations had the same number of analysts. Thus, hierarchies had more DMUs than did teams (13 as opposed to 9). Having extra personnel, however, did not enhance the hierarchies' performance. Consider for a moment a team with 13 analysts. In such a team, each analyst faces a less complex subproblem and the degree of information loss when the team votes to form the final decision is less (a reduction from 27 to 13 rather than to 9). Consequently, a team with the same number of DMUs as the hierarchy would perform even better than a team with the same number of analysts. The benefit of these "extra decision makers" to the hierarchy is that they store a reduced set of information. What this analysis suggests, however, is that such storage does not, under many conditions, make up for the loss in information that the organization suffers from not treating all decision makers as peers (as is done in the team). This is particularly true when the task is unbiased.

In the proposed model, only analysts leave and all analysts are equally likely to leave. Yet, in real organizations, upper level management does change and there are often systemic controls on who leaves. For the hierarchy the retention of upper level management results in greater institutional memory; consequently, were management to leave, the hierarchy might have suffered more from turnover than it did. For both teams and hierarchies were leaving the organization based on tenure or on poor performance slightly different results would have followed. For the task examined performance and experience are correlated. Thus firing low performing individuals is like firing new workers and higher turnover rates would be no more effective than lower rates; whereas, firing individuals with more tenure is like firing high performers and should decrease final performance. Since individuals with more tenure have more experience making all individuals equally likely to leave is effectively the same as following a mixed firing strategy.

Turnover leads to the loss of information, a reduction in institutional memory, and generally lower performance in the proposed model as there is no repository for knowledge in the organization other than personnel. External repositories for knowledge, such as rules of operation, forms in file cabinets, and computerized data bases (like those that occur in information, decision, or expert systems) are not taken into account. Yet, such repositories as the trappings of bureaucracies [17] affect organizational behavior [36]. Were the model expanded to consider such external repositories the organization might be less affected by turnover and would probably exhibit a greater rate and level of learning when novices or "good fit" personnel are hired. Whether such repositories are beneficial to the organization when inappropriate personnel are hired is more problematic.

In the proposed model redundancy is modeled by increasing the overlap in information across decision makers in a rotating fashion. Thus as redundancy increases so does the complexity of the subproblem seen by the analyst but not the degree of consensus across decision makers. A main effect of redundancy is therefore to cause the organizations to learn not only less but slower. Alternative models of redundancy would have lead to different results. For example, if redundancy was modeled by multiple analysts having completely identical information then redundancy would have consistently mitigated the effect of turnover, there may have been an optimum level of redundancy per turnover rate, but in the absence of turnover the organization would still have learned slower the greater the level of redundancy.

Despite the limitations of the model it does capture many of the features of organizational behavior and thus can serve as a framework to look at many of the issues just discussed. For example, using the proposed model as a base different models of turnover and redundancy could be examined. As another example, the approach used herein could be easily adapted to explore whether the experience possessed by the new personnel affect performance, or whether different structures - such as matrix - perform better under these various conditions.

5. Coordinating for Success

An analytical exploration was conducted of the relationship between the organization's coordination scheme and organizational learning when the organization is faced with a quasi-repetitive integrated decision making task and there is personnel turnover. This exploration suggests that the coordination scheme - the organizational structure and the information access structure - impact not only performance but the rate at which the organization learns and the ability of the organization to cope with turnover.

No coordination scheme, however, dominates. Teams outperform hierarchies when the task is unbiased but hierarchies can outperform teams when the task is biased if they have the right information access structure and right level of turnover. Thus, there is at times value to the organization in having different types of agents (such as in the hierarchy). Different coordination schemes are optimal for different tasks. Choosing the wrong scheme can actually damage performance. For example, if you think the task is unbiased and you expect only minimal turnover then you would be best off to organize as a team with no redundancy. If you are wrong, and the task is biased you would have been better off to have some redundancy or to have organized as a hierarchy. For both hierarchies and teams, however, having a little redundancy may enable the organization to operate at an acceptable, although not optimum, performance level thus making the organization somewhat immune to needing to know task type.

Different information access structures produce different organizational behavior as when information redundancy increases at least three different forces come into play - complexity, resolution, and overlap. Complexity occurs as each decision maker has more information to contend with, there are more classes of subproblems, the frequency of each subproblem is less, and so the individual, and consequently the organization, learns more slowly. For the decision maker, receiving more information on a specific task decreases his or her ability to notice similarities between tasks. While this may lead to greater eventual learning, it will slow down the rate of learning. Resolution occurs as each decision maker by having access to more information has a greater chance of estimating the true answer. Overlap occurs as pieces of information are seen by multiple decision makers. It is this effect that is typically associated with redundancy and forms our expectation that redundancy should mitigate the impact of turnover.

This study suggests, however, that this expectation may in many cases be wrong. The explanation lies in the fact that although increased redundancy may lead to increased overlap in who has access to what it may not increase consensus as different decision makers by sharing some and not all information may come to interpret the information they share in different ways. Indeed, the indiscriminate distribution of information across decision makers in order to assure multiple access can degrade performance and will often slow the rate of learning. The tradeoff between task complexity and redundancy that this analysis illustrates suggests that a *need to know strategy* may actually improve organizational performance by reducing task complexity. In any case, this analysis suggest that the manager must be careful in assigning access in order not to reduce performance by making tasks too complex. Thus overall, the best coordination scheme may be one of divide and conquer.

REFERENCES

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

[2] March J.G. and J.P. Olsen, "The uncertainty of the past: organizational learning under ambiguity", *European Journal of Political Research*, Vol. 3, 1975, pp. 147-71.

[3] Padgett J.F., "Managing Garbage Can Hierarchies", *Administrative Science Quarterly*, Vol. 25, No. 4, December, 1980, pp. 583-604.

[4] March J. and R. Weissinger-Baylon (Eds.), *Ambiguity and Command: Organizational Perspectives on Military Decision Making*, Pitman, Boston, MA, 1986.

[5] Simon H.A., *Administrative Behavior*, Macmillan, New York, NY, 1947.

[6] March J.G. and H.A. Simon, *Organizations*, Wiley, New York, NY, 1958.

[7] Cyert R.M. and J.G. March, *A Behavioral Theory of the Firm*, Prentice-Hall, Englewood Cliffs, NJ, 1963.

[8] Steinbruner J.D., *The Cybernetic Theory of Decision Processes*, Princeton University Press, Princeton, NJ, 1974.

[9] Durfee E., Lesser V. and D. Corkill, "Increasing coherence in distributed problem solving networks", *Proceedings of the Ninth International Conference on Artificial Intelligence*, 1985.

[10] Lesser V.R. and D. D. Corkill, "Functionally accurate, cooperative distributed systems", *IEEE Transactions on Man, Systems, and Cybernetics*, Vol. SMC-11, No. 1, 1981, pp. 81-96.

[11] Reid D., "An Algorithm for Tracking Multiple Targets", *IEEE Trans. Automatic Control*, Vol. 6, No. AC24, 1979, pp. 843-854.

[12] Smith R., "The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver", *IEEE Trans. on Computers*, Vol. C-29, No. 12, 1980, pp. 1104-1113.

[13] Steeb R., et. al, "Distributed intelligence for air fleet control", Tech. report, The Rand Corporation, 1980.

[14] Thorndyke P., McArthur D. and S. Cammarata, "Autopilot: A distributed planner for air fleet control", *Proceedings of the Seventh IJCAI*, Vancouver, B.C. Canada, 1981, pp. 171-177.

[15] Davis R., "Report on the Workshop on Distributed Artificial Intelligence", *SIGART Newsletter*, Vol. 73, 1980, pp. 43-52.

[16] Davis R., "Report on the Second Workshop on Distributed Artificial Intelligence", *SIGART Newsletter*, Vol. 80, 1982, pp. 13-23.

[17] Weber M., *Bureaucracy*, Oxford University Press, Oxford, England, 1922.

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

[19] Mintzberg H., *The Structure of Organizations*, Prentice Hall, Inc., NJ, 1979.

[20] Mintzberg H., *Structure in Fives: Designing Effective Organizations*, Prentice Hall, Inc., NJ, 1983.

[21] Burns T. and G. Stalker, *The Management of Innovation*, Tavistock, London, England, 1961.

[22] Carley K.M., *Measuring Efficiency in a Garbage Can Hierarchy*, Pitman, Boston, MA, 1986, ch. 8.

[23] Marschak J., "Elements for a Theory of Teams", *Management Science*, Vol. 1, 1955, pp. 127-137.

[24] Arrow K.J. and R. Radner, "Allocation of Resources in Large Teams", *Econometrica*, Vol. 47, 1979, pp. 361-85.

[25] Radner R., *Decentralization and Incentives*, University of Minnesota Press, Minneapolis, MN, 1987.

[26] Gloves T. and J. Ledyard, "Optimal Allocations of Public Goods: A Solution to the Free-Rider Problem", *Econometrica*, Vol. 45, 1977, pp. 738-809.

[27] Tsitsiklis J.N. and M. Athans, "On the Complexity of Decentralized Decision Making and Detection Problems", *IEEE Transaction on Automatic Control*, Vol. AC-30, 1984, pp. 440-446.

[28] Strand R.G., "An Efficient Suboptimal Decision Procedure for Associating Sensor Data with Stored Tracks in Real-Time Surveillance Systems", *Proc. IEEE Conf. on Decision and Control*, Miami Beach, FL, 1971, pp. 33-37.

[29] Bar-Shalom Y. and E. Tse, "Tracking in a Cluttered Environment with Probabilistic Data Association", *Proc. of the Fourth Symposium on Nonlinear Estimation*, University of California, San Diego, CA, 1973, pp. 13-22.

[30] McLain D.H., "Drawing Contours from Arbitrary Data Points", *The Computer Journal*, Vol. 17, 1974, pp. 318-324.

[31] Wilkinson L., *SYGRAPH*, SYSTAT, INC., Evanston, IL, 1988.

[32] Tversky, A. and D. Kahneman, "The framing of decisions and the rationality of choice", *Science*, Vol. 211, 1981, pp. 453-458.

[33] Tversky, A. and D. Khaneman, "Availability: A heuristic for judging frequency and probability", *Cognitive Psychology*, Vol. 5, 1973, pp. 207-232.

[34] Nisbett R. and L. Ross, *Human inference: strategies and shortcomings of social judgment*, Prentice-Hall, Englewood Cliffs, NJ, 1980.

[35] Argote L., Beckman S. and D. Epple, "The persistence and transfer of learning in industrial settings", *Paper presented at the St.Louis meetings of the Institute of Management Sciences (TIMS) and the Operations Research Society (ORSA)*, Vol. , 1987, pp. .

[36] Levitt B. and J.G. March, "Organizational Learning", *Annual Review of Sociology*, Vol. 14, 1988, pp. 319-40.

[37] March J.G. and Z. Shapira, *Behavioral decision theory and organizational decision theory*, Kent Publishing Co., Boston, MA, 1982.