CHAPTER

20

Organizational Performance, Coordination, and Cognition

Kathleen M. Carley
Carnegie Mellon University

Organizations can be viewed as collections of intelligent agents who are cognitively restricted, task oriented, and socially situated. Accordingly, the behavior of the organization is affected by the behavior of the agents and by how the agents are coordinated. Coordination can be achieved passively through the extant organizational structure. This structure limits agent behavior by determining who has access to what information, who must make which decisions, and who must report what to whom.

Carley and Prietula (1994) referred to this perspective as ACTS theory and described it in detail. Central to this perspective is the idea that organizational performance is jointly affected by coordination and cognition. The thesis is that organizational performance should change, given a particular coordination structure, as you replace the agents with agents of differing cognitive abilities. Replace people by robots or rocks and the organizational performance should change. Similarly, organizational performance should change, given a particular type of agent, as you alter the coordination structure. Change the organizational structure from an democratic team to a more hierarchical structure and organizational performance should change. Similar arguments have been made at the interorganizational level (Malone, 1986; Williamson, 1975).

These arguments seem obvious, yet much of organization theory has looked at organizational performance as being dependent on the coordination structure sans agent cognition. For example, structuralism, institutionalism, and population ecology suggest that organizational performance is

Kathleen M. Carley, 2001, "Organizational Performance, Coordination, and Cognition." Ch. 20 (pp. 595-622) in Olson, G.M., Malone, T.W., & Smith, J.B. (Eds) *Coordination Theory and Collaboration Technology*. Mahwah, NJ: Lawrence Erlbaum Associates.

largely determined by factors other than human cognitive and affective behavior. In contrast, much of the research in organizational behavior has focused on the impact of cognition sans coordination. We have only a limited understanding of the organization as a collection of coordinated intelligent agents. Recent advances have been made, however, using meso-level models in which macrolevel organizational behavior emerges from the microlevel agent actions (Carley, 1991, 1992; Carroll, 1984; Masuch & LaPotin, 1989). In these models, organizations are formed as collections of intelligent adaptive agents. The models of the agents are based on reasonable assumptions about human behavior, generally predicated on decades of research. The complexity of these agents is sufficient that recognizable, and important, organizational behaviors emerge. In this way, these models have increased our understanding of organizational performance and demonstrated the value of meso-level adaptive agent modeling for building organizational theory. All this being said, there has been little attention as to whether the cognitive nature of the agent alters organizational performance and whether there is an interaction effect between cognition and coordination.

Clearly, organizations composed of agents with different features behave differently. Cohen, March, and Olsen (1972) found that the amount of effort agents expended affected the quality of organizational decisions. Carroll (1984) found that organizations composed of agents with different cultural biases perform differently. Carley, Park, and Prietula (1993) found that whether or not agents lied affected the degree to which the organization wasted time. Lin and Carley (1993) found that organizations of proactive agents tended to outperform organizations of reactive agents. Numerous other examples exist. In all these analyses, and many others, we find that agent features make important differences in organizational performance. However, none of these analyses indicates whether or not increasing the cognitive realism of the agent models, so that they more closely approximate the human agent, alters organizational performance.

Organizations employing different coordination structures may also vary in their performance. Mackenzie (1978) and Roberts (1989) argued that hierarchy is linked organizational efficiency and reliability. Galbraith (1973, 1977) discussed the relative importance of centralization and decentralization. Thompson (1967), Mintzberg (1979), La Porte and Consolini (1991), and Roberts (1990), argued that loosely coupled or structural redundant organizations are high performers in stressful conditions. Numerous other examples exist. In all these analyses, and many others, we find that organizations who coordinate through the use of different designs exhibit different performance. Collectively these studies demonstrate that there is no one best organizational design. However, none of these analyses indicate whether or not the performance attributable to a particular coordination structure will remain constant as the agent model is altered.

Understanding whether the degree of agent veridicality interacts with structural changes in the organization or the task in important ways is important for advancing organizational theory. If we find that the veridicality of the agent model does not interact with the organizational structure or task, that the performance of the coordination structures and task are constant regardless of the agent model, then macro theoretical approaches that ignore the agent gain support. In contrast, if we find that organizational performance, and particularly the relative performance of the different coordination structures, is dependent on the realism of the agent model then these theoretical approaches are called into question.

In this chapter we address this issue directly by contrasting the performance of organizations with different coordination structures, different task complexities (from the agent's perspective), and composed of different "types" of agents. The performance of organizations composed of simple adaptive agents (ELM agents), complex adaptive agents (Soar agents), and humans is examined. As we move from ELM to Soar to humans, presumably the realism of agents is increasing. At issue is how this realism interacts with organizational and task constraints in affecting organizational performance. Similarly, the performance of organizations with teams and hierarchies, blocked and distributed information are examined. For each organization performance is measured. Using this data the relative impact of, and interactions among, cognition and coordination on organizational performance is explored relative to a simple ternary classification–choice task.

THE TASK

The task faced by each organization, regardless of its coordination scheme or the cognitive architecture of its agents, is a ternary classification–choice task. Without loss of generality we can think of this task as a highly stylized radar task. There is a range of physical air space surrounding the radar equipment. This airspace can be scanned by the radar equipment and information about the flying object can be gathered. Within the airspace there is a single object. This object has a true state that is either: FRIENDLY; NEUTRAL; or HOSTILE. This object has nine features. These are: speed (mph); direction (indicating degrees of deflection by which the flight path deviates from a direct route); range (miles); altitude; angle; corridor status (in, edge, out); identification (friendly, civilian, unknown); size (feet; small, medium, large); radar emission type (weather, none, weapons; Carley & Lin, 1992). Each feature has a value of either: 1, 2, or 3. The interpretation of these values depends on the feature (e.g., if the feature is "speed," the Value 1 means speed is low, Value 2 means speed is medium, and Value 3 means speed is high). Initially agents in the organization do not know whether having a low or high value on one feature or another is associated with the object being truly FRIENDLY, NEUTRAL, or HOSTILE. Agents must learn these associations.

The true status of the aircraft is defined external to the organization and is not manipulatable by the organization. This is the characteristic of the design for the task. The true status of the aircraft is manipulated by the experimental designer. By changing the rule relating a pattern of aircraft characteristics to an outcome the researcher can examine different types of tasks. In this chapter, the true status of aircraft is generated by using decomposable and unbiased scheme for defining the task.[1]

In an unbiased environment all three possible outcomes are equally likely. In a decomposable environment each piece of information is equally important. Consequently no analyst plays a more important role than any other simply on the basis of seeing a certain type of information. The task has a complexity level of nine; that is, there are nine pieces of information F1 through F9. Each piece of information can take on one of three values, FRIENDLY = 1, NEUTRAL = 2, and HOSTILE = 3. The true state of the aircraft is defined on the basis of the sum of these nine features. If the sum of the values for these features for a specific aircraft is less than 17 the true state is friendly, if this sum is greater than 19 the true state is hostile, otherwise the true state is neutral. This defining rule establishes which true state is associated with which pattern of information. The members of the organization do not know apriori how the true state is calculated from the set of features. Consequently, the agents in the organization do not know apriori how to relate a particular pattern of information to a particular outcome. Because there are nine pieces of information each of which can take on three values there are 19,683 possible patterns that the organization needs to learn. The organizational design limits the organization's ability to learn these patterns.

Each analyst must decide which one of three states (FRIENDLY, NEUTRAL, HOSTILE) the passing object is, based on the information (the features of the object) he or she can access. The information known by each agent is a subset of the total information, and how many pieces of information in each subset is dependent on the organization's resource access structure. After seeing the information, each agent has to make a decision and deliver a recommendation. How these recommendations are processed

[1]In a decomposable task each component has a separable, identifiable, and additive effect in determining the problem solution. Each piece of information contributes equally to the final decision. No agent has greater "power" simply by virtue of having access to a more powerful or more important piece of information. In an unbiased environment approximately one third of the 19683 aircraft are friendly and one third of the aircraft are hostile. This is an environment where all the possible outcomes are equally likely to be true (Carley & Lin 1992).

or combined by the organization depends on the organizational structure. After each analyst has made its recommendation it receives feedback on the true state of the aircraft. The feedback can be considered as part of the training procedure.

This radar task is based on a real-world problem, and variations of it have been widely examined (Carley, 1990, 1991, 1992; Hollenbeck, Sego, Ilgen, & Major, 1991; Mallubhatla, Pattipati, Tang, & Kleinman, 1991). Two features of this task make it appropriate for our present purpose. First, the true state of the object is known. Thus feedback can be provided and issues of training (and hence differences in learning procedures) can be addressed. Second, this task is complex enough that it can be solved in a distributed environment where information is shared by different agents, and multiple agents can be used to work on different aspect of the task.

MODELS OF COORDINATION

The organizational design can serve as a passive coordination scheme defining who does what when. Herein, two aspects of organizational design are considered: the organizational structure and the resource access structure. The organizational structure defines who reports to whom and how the organization makes its decision. The resource access structure defines who has access to what information or resources. Regardless of the organizational design there are nine analysts. Each analyst, regardless of the organization it finds itself in, has access to three pieces of information on each task, makes a recommendation based on this information, and passes on this recommendations as its decision. For each analyst each of the three pieces of information can take on three different values. If there is a manager in the organization then the manager takes these nine recommendations and makes a decision based on them. If there is no manager the organizational decision is simply the majority vote. In principle, each manager has the possibility of seeing nine different pieces of information each of which can take on three different values.

Two organizational structures are examined (see Fig. 20.1): the team with manager and the team with voting. Within the team with manager structure, the organizational decision is made by the CEO who makes this decision after it receives all nine analysts' decisions. Within the team with voting structure, the organizational decision is the majority decision given the separate decisions provided by the nine analysts. We focus on team structures as previous research has demonstrated that teams learn more quickly and are typically more "disturbed" by any type of internal or external stress such as turnover or missing information than are other organizational structures (Carley, 1990, 1991, 1992). The two team structures examined differ, however, in their centralization. The team with voting is a decentralized structure, whereas the team with manager is a centralized structure. By focusing on these structures we will be able to see whether slight changes in the cognitive makeup of the agents make major differences in organizational performance when the decisions made by these agents are combined in different ways. Such an analysis would be more difficult with more complex organizational structures although later work should examine such structures.

Two resource access structures are examined (see Fig. 20.1): distributed and blocked. Within the blocked resource access structure, multiple agents see exactly the same information. In this case, three analysts see the same three pieces of information on each task. Within the distributed resource access structure, no two agents see exactly the same information and each piece of information is seen by more than one person. Specifically, each piece of information is seen by three different analysts. In both resource access structures each analyst sees the same number of pieces of information (three). Thus, regardless of the structure, all analysts are facing tasks with the same complexity and so their information load is the same. All managers, regardless of the resource access structure, see nine pieces of information and so are facing the same information load. Differences in "learning" at the analyst level can only be attributable to differences in the cognitive realism of the agents and not to differences in information loads. Differences

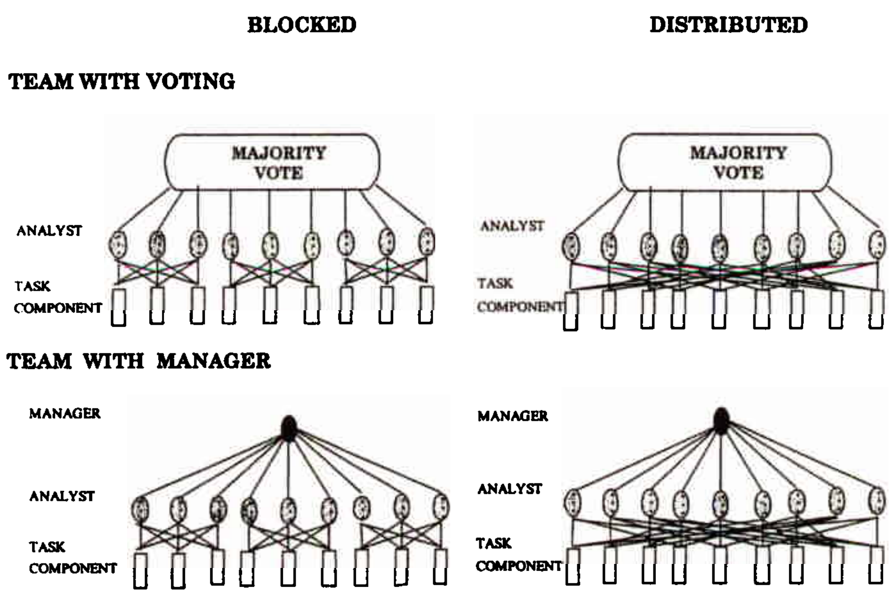

FIG. 20.1. Organizational coordination schemes.

in learning at the managerial level can only be attributable to the extent to which the information they have is consistent and to the cognitive realism of the agents. By focusing on these resource access structures we will be able to see whether slight changes in the cognitive makeup of the agents exacerbate differences due to the level of similarity in different agents' mental models.

For analysts, since they see three pieces of information each of which can take on three values, there are 27 possible patterns. For managers, since they see nine pieces of information each of which can take on three values there are 39 or 19,683 possible patterns. Not all of the patterns are equally likely to occur. The likelihood of a specific pattern is a function of the set of tasks that the organization sees and, for managers, it is also a function of the resource access structure. In this study, the tasks were chosen randomly such that all patterns are equally likely. Thus, for analysts, all patterns are equally likely. For managers, however, the resource access structure determines the likelihood of specific patterns. If the agents are in an organization with a blocked resource access structure then the manager may potentially see only 27 of the possible 19,683 patterns; however, the number actually seen depends on the type of agent. Whereas, a manager in an organization with a distributed resource access structure has a greater likelihood of seeing all 19,683 patterns (if the organization faces that many tasks). How many patterns the manager actually sees depends on the model of cognition. Whether these differences in the number of possible patterns, the probability that certain patterns will occur, and the number of potentially observable possible patterns will affect learning or the speed of making a decision depends on the agent's cognitive capabilities. Regardless of the agent's cognitive makeup the number of patterns affects the resolution of the information available to the decisionmaker and so may affect the quality of the decision maker's decision. For the environment being examined there are 19,683 unique aircraft; hence, 19,683 patterns relating the aircrafts nine features to its true state. The fewer patterns an agent can see the less resolution the agent has on the overall problem.

MODELS OF COGNITION

Three different agent "models" are considered: ELM agents, Soar agents, and human agents. Both ELM (Carley, 1991, 1992; Carley & Lin, 1992) and Soar (Papageorgiou & Carley, 1992; Carley, Kjaer-Hansen, Prietula, & Newell, 1992; Carley, Park, & Prietula, 1993) agents have been used in models of organizations composed of artificial adaptive agents; however, their performance has not been contrasted. Further, there is no research demonstrating how the behavior of these artificial agents actually compare with the behavior of human agents given the specific task we employ. ELM agents are based on a simple experiential learning model using an incremental adaptive algorithm similar in intent to those used in classic learning theory (Bush & Mosteller, 1955). Soar agents are knowledge intensive agents who are capable of employing the various common search algorithms for problem solving (Newell, Yost, Laird, Rosenbloom, & Altmann, 1991).[2] The important point here is that ELM agents are task specific in nature and employ a learning and decision procedure that reflects only a little of what is known about human cognition. In contrast, Soar agents are not task specific and the learning and decision procedures in Soar have been shown to be consistent with much of what is known about human cognition (Laird, Newell, & Rosenbloom, 1987; Rosenbloom, Laird, Newell, & McCarl, 1989). Clearly Soar agents are much closer to humans than are ELM agents in terms of individual behavior. We ask, does this closeness matter when it comes to examining organizational behavior.

Both the specific ELM (Carley & Lin, 1997; Lin & Carley, 1977) and Radar-Soar (Ye & Carley, forthcoming) agents that we employ have been described in detail elsewhere.[3] Soar itself has been described in numerous reports (Laird, Newell, & Rosenbloom, 1987; Laird, Rosenbloom, & Newell, 1986). We limit our description of these models to a brief overview of how these models work, mainly to highlight differences in their assumptions about human problem solving behavior as it relates to the ternary task used in this analysis.

Regardless of how the agents are modeled they make their decisions on the basis of the same incoming information. How the agents are modeled affects how they access, use, recall, and make decisions on the basis of this information. In this study, all analysts see three pieces of information which are the raw task information (about the aircraft). If the agent is a manager he or she sees nine pieces of information which are the decisions of the nine analysts. Each piece of information can take on one of three values: FRIENDLY (=1), NEUTRAL (=2), and HOSTILE (=3). The values for the set of information seen by the agent is referred to as a pattern. For example, an agent might see the pattern 111, meaning that the agent sees three pieces of information for which all three have a value of one. What specific pieces of information the agent sees depends on his or her position in the organization's coordination scheme; specifically, in the organizational structure and the resource access structure.

Regardless of how the agents are modeled they all receive the same feedback in the same way. During training, after each agent has reported his or her decision, he or she is told what the true answer is for the entire problem. That

[2]For general discussions of Soar see Mitchell (1988a, 1988b).

[3]See also Carley (1992) for a description of the ELM agents when faced with a binary task. See also Carley and Newell (1994) for a detailed description of Soar as a model of the human agent.

is, each agent is told the true state of the aircraft. Agents are not told how well the organization or they themselves are doing. They have to infer that information from the feedback they receive about the aircraft's true state.

ELM Agents

The Experiential Learning Model (ELM) is a simple model of agent learning and decision making in which the agent incrementally adapts its behavior on the basis of feedback. This model is similar in kind to Bush and Mosteller's (1955) type learning models. This model has been used within various organizational simulation test beds, such as CORP (Carley & Lin, 1997; Lin & Carley, 1977). Such testbeds have been used to examine the relative performance of organizations which place different sociocultural–historical constraints on the agents. ELM employs a simple situated-cognition model of individual action in which agents learn from experience. What information agents have access to and what action they can take are dependent on their social situation in the organization and specific task assigned to them. In ELM the social situation is defined by the agent's position in the organizational design (Carley & Prietula, 1994).

ELM agents are boundedly rational in the sense that: (a) they are not omniscient, (b) they can act only on the information available and their historical knowledge; (c) their historical knowledge is limited to information about the distribution of previous events; and (d) their knowledge is task specific and there are no built in procedures for transferring knowledge between tasks. ELM agents are adaptive in the sense that they change their memories over time.

ELM agents' memories can be thought of as a series of rules of the form "if pattern *x* is seen then report *y*." Each agent has an information load equal to the number of possible patterns that it can possibly see. This information load is affected by the organizational design. In the designs examined the information load for all analysts is 27 and for all managers is 19,683. The higher the load the slower the agent loads and the longer is takes to make a decision. It does not necessarily affect the accuracy of the agent's decision.

The agent reports whether he or she thinks that the aircraft is FRIENDLY (=1), NEUTRAL (=2), or HOSTILE (=3). Each ELM agent has in its memory information for each pattern as to the number of times that he or she saw this pattern and the true answer was FRIENDLY, NEUTRAL, or HOSTILE. The agent then selects as its decision for an observed pattern that event that was most likely in the past. Thus, when faced with the pattern 111, the agent will recall the distribution of true states associated with this pattern, that is, FRIENDLY occurred 6 times, NEUTRAL occurred 4 times, and HOSTILE occurred 2 times. Then the agent will provide as its answer the most likely event, in this case, FRIENDLY.

During training, while the agent is learning, the agent builds up these distributions incrementally. Thus, agents who observe different sets of patterns will actually learn different behavior. If all patterns are equally likely then the analysts and managers will, in the limit, learn to act approximately as majority classifiers (given the task being examined).

For a specific number of tasks, managers in a blocked resource access structure, because they see at most 27 of the possible patterns, will build up more information on how those 27 patterns relate to the three possible outcomes. For those same tasks, managers in a distributed resource access structure have the potential to see 19,683 and so are less likely to build up information on any one pattern. Whether having more information on fewer patterns or less information on more patterns will lead to better decisions is the issue at the managerial level when comparing these resource access schemes for the ELM agents.

Radar–Soar Agents

Radar–Soar agents are complex adaptive agents built on top of the Soar architecture. Following is a brief description of Soar. This is followed by a brief description of the Soar agents.

Soar. Soar is a general-purpose program for solving problems. It incorporates specific knowledge about the world as a set of rules that guide it in solving problems. Soar agents learn from experience by remembering how they solve problems—this is referred to as chunking. Within Soar all cognitive behavior is considered to be symbolic and goal oriented. Soar is considered to be unified theory of cognition as it is a single, integrated set of information processing mechanisms that try to explain every aspect of human thought, not one or two experimental results (Newell, 1990).

Soar characterizes all cognitive behavior as search in problem spaces and serves as an architecture for general intelligent behavior (Laird et al., 1987). Soar's structure is built in levels, starting with memory and proceeding to decisions and goals. A learning procedure (chunking) and default knowledge are also incorporated into the system, but need not be used (Rosenbloom et al., 1989). In this chapter the chunking mechanism is not employed.

• *Memory:* All Soar's long-term knowledge is stored in a single production memory composed of if–then rules. Memory access consists of the execution of these productions. As these productions are executed information is retrieved into a global working memory. The working memory is a temporary memory which roughly corresponds to the set of things that the agent is attending to at any given moment. A special type of working memory structure is the preference. Preferences encode control knowledge about the acceptability and desirability of actions. Acceptability preferences determine which action should be considered as a candidate action. Desirability preferences define a partial ordering on the candidate actions.

• *Decision:* The decision cycle requires two phases: elaboration and decision. During the elaboration phase, the long-term (production) memory is accessed repeatedly (effectively in parallel), until no more productions can execute. This can result in a set of preferences being established. During the decision phase, these preferences are interpreted. This can result in changes in the agent's goal, state, actions, etc. This decision cycle ensures that Soar will make its decisions after all the rules have been fired; consequently, Soar will use the most powerful knowledge it has available. When there is little knowledge in long-term memory, Soar will behave in ways that resemble general problem-solving techniques such as hill climbing, or means–ends analysis. When there is lots of knowledge in long-term memory, Soar will behave as an expert as it will have clear preferences about what to do next.

• *Goals:* Goals are set whenever a decision cannot be made; that is, when an impasse is reached during the decision phase. Impasses include: ties; conflicts; no-changes; constraint failures. When an impasse occurs, Soar creates a subgoal to resolve the impasse and a corresponding performance context. This results in a hierarchical goal structure. A subgoal is terminated when either its impasse is resolved or a higher impasse in the stack is resolved. This architectural feature is called *universal subgoaling* (Laird et al., 1986). Goals are functions on behavior (i.e., agents prefer some actions to others). An agent's behavior is determined by the *principle of rationality*; that is, if the agent knows that one of its actions leads to a preferred situation according to its goal, then it will intend the preferred action, and this action will then occur if it is possible (Newell et al., 1991). Agents exhibit goal-directed behavior, because all actions intend to attain the agent's goal. Agent's are rational, because everything the agent knows serves the agent's interest.

Radar–Soar

Developing task-related agents in Soar requires determining for each agent its initial task knowledge, possible actions, and problem spaces. This task knowledge is instantiated as productions in the agent's long-term knowledge base. The actions are instantiated as operators within problem spaces (which are the arenas for action). As productions are fired the agents solve problems by moving through a series of problem spaces and within each space taking those actions that are preferred. A task is formulated using a problem space by: determining which problem space to adopt; setting a goal which determines which desired state is adopted; and determining the initial state. The formulated task is accomplished by: attempting state-operator pairs; applying operators; and terminating the current state when the desired state is reached.

Analysts and managers differ in the problem spaces and actions they used for communication. Both analysts and managers move between problem spaces by taking actions in one space which then move the agent to a subsequent space. Regardless of whether the agent is an analyst or a manager the agent makes its decision within a make-decision space.

The make-decision space is critical to the subject of this chapter. In this space the agent compares the newly observed information with each of the models currently in its knowledge base and calculates the level of match for each model. A model is a description of aircraft features and the values they take on (a pattern) and a predicted outcome. Each model is represented as a rule of the form "if pattern X is seen then report Y." (These patterns are equivalent to the patterns in ELM.) The match is the number of features in the observed aircraft and the model that have identical values. For example, for analyst, if the model is that all three features are FRIENDLY and the observation is that one of the three features is FRIENDLY and the other two are not, then the match is one. If the agent has *N* models then *N* matches are calculated. The agents preference for a model is based on the match. Specifically, agents prefer models with higher matches and are indifferent among those with the same level of match. After the matches have been calculated the agent chooses that model that has the highest match. If there are several such models the agent randomly chooses among them. Among the equally preferred models, which model is chosen is determined randomly. Because the agent stores a model for each problem that it observes, this results in the agent's decision having a probability associated with it proportional to the number of times the agent has observed this type of aircraft with this type of outcome. The more models the agent has the longer it takes for the agent to make a decision (without chunking). When a model is chosen the agent makes as its decision the choice recommended by this model.

When the current situation is beyond the agent's knowledge, he or she will make a decision based on reasoning and not simply by guessing. If an agent sees an aircraft he or she will compare the features of that aircraft to the models available (to start with this will only be the initial knowledge). Then the agent will choose the model that has the closest match to the current aircraft. For example, imagine that the agent sees only three aircraft features and that the first aircraft seen has the features high, high, medium. The agent will calculate the match with the existing models. The match with the first model (all high) is 2, the match with the second model (all medium) is 1, and the match with the third model (all low) is 0. These matches set up a preference ordering among outcomes such that the first model is preferred to the second, which is preferred to the third. Thus the agent will choose the first model and make whatever decision it suggests.

A second critical space for this model is the update knowledge space. The update knowledge space is the space where the agent creates new models. In this space the agent takes its observation and the feedback it has received and creates a new model. Feedback is of the form the aircraft's true status is FRIENDLY (or NEUTRAL, or HOSTILE). Each observation results in a new model. Each time the agent receives feedback during training it creates a new model linking the observed pattern with the true state of the aircraft.

• *Analysts:* Overall, the analyst's goal is to resolve all commands forwarded by the manager. The analyst sees a sequence of commands and responds to these, and stops when all commands are resolved. These commands direct the analyst in the observation of the airspace and the making of decisions on the information observed. Each Radar–Soar analyst is implemented with 11 problem spaces connected hierarchically. The top problem space contains the initial states and the desired state (the goal state). The goal state is where all commands are resolved. This connects to the communication problem space. In the communication problem space, three operators are sequentially proposed: the get-command operator; the report operator; and the get-feedback operator. Each operator in turn calls the corresponding problem space. Within the get-command problem space, there are two kinds of commands from the manager: "observe" and "tell-me," which when accepted by the analyst lead to two different problem spaces, the observe air-space problem space and the make-decision problem space. In the observe air-space problem space, two operators are raised. One is the parse information operator which the analyst uses to analyze the signal captured from the air-space about the flying object; another is the interpret operator which the analyst uses to convert each signal to the attribute it represents. Each operator calls up their corresponding problem space. In the make-decision problem space, two operators are also proposed, one is the model-select operator which compares the information of the aircraft to all the different models in an analyst's long-term memory, then makes a decision based on the maximum match; the other operator is the write-decision operator which the analyst uses to record his or her decision so that it can be reported to the manager later. Each of these operators also creates their own problem spaces in which the operators are enacted.

Each analyst sees exactly three pieces of information or features. The three initial productions (or initial models) are:

If feature-1 = feature-2 = feature-3 = FRIENDLY then decision = FRIENDLY
If feature-1 = feature-2 = feature-3 = NEUTRAL then decision = NEUTRAL
If feature-1 = feature-2 = feature-3 = HOSTILE then decision = HOSTILE

Within the make-decision space the analyst's goal is to suggest a decision about the true state of the aircraft. Initially the analyst can only compare the observation with these three models. Thus, initially the analyst will act as a majority classifier. Over time as new models are built the analyst's behavior will come to emulate the lessons of history. However, to the extent that multiple models are equally valid the analyst will choose between them stochastically.

• *Managers:* Managers are conceptually similar to analysts. The main difference is that they have additional communication actions, can command other agents, and their initial models are based on nine rather than three pieces of information. The Radar–Soar manager is implemented using six problem spaces connected hierarchically. The top problem space contains the initial states and the desired state for the manager. The top problem space, as with the analyst connects to the communication problem space. Within the communication problem space four operators are sequentially proposed: the give-command operator; the receive-response operator; the make-organizational decision operator; and get-feedback operator. These four operators in turn lead to four problem spaces in which the manager carries out its actions.

The manager's goal is to make the best possible organizational decision for each of the problems it faces. Each manager sees exactly nine pieces of information, one for each analyst. The three initial productions are:

If agent-1 = agent-2 =. . . agent-9 = FRIENDLY then organizational decision = FRIENDLY

If agent-1 = agent-2 =. . . agent-9 = NEUTRAL then organizational decision = NEUTRAL

If agent-1 = agent-2 =. . . agent-9 = HOSTILE then organizational decision = HOSTILE

Ed: no space before ellipsis (as on ms) ok /2

The manager will build up models at the same rate as the analysts. A manager in a blocked resource access structure will build up more models with similar left hand sides than will managers in distributed resource access structures. Hence, when a new aircraft is observed there will, on average, be more models with equivalent matches to the new data. Hence the manager's preference ordering will be less likely to suggest a unique outcome and the manager will be more likely to be indifferent among a wider variety of models. Consequently, stochastic factors may play a greater role in the organizational outcome when the resource access structure is blocked.

Because Radar–Soar agents act in a stochastic fashion, two agents who see the same information may respond differently. Thus, in a blocked resource access structure, even though three analysts see the same information they may not respond in a similar fashion. Consequently, the manager in a Radar–Soar organization, even when the resource access structure is blocked, may potentially see more patterns than the ELM manager in the same situation. Whether the potential lack of consensus among analysts despite common incoming information, in the Radar–Soar situation, will lead to better decisions is the issue when comparing Radar–Soar and ELM organizations with blocked resource access structures. Whether the lack of consensus among analysts caused by interpretation (differences caused by analysts in a blocked resource access structure stochastically making different choices) or caused by different information (differences caused by analysts in a distributed resource access structure making different choices) leads to better or worse performance is the issue in contrasting Radar–Soar organizations under the different resource access structures.

Human Agents

A series of experiments were run by Carley and Prietula to examine the effect of structure. These organizations duplicate the organizations examined via simulation by Carley.[4] In the human experiments subjects were run as either analysts or managers; managers and analysts are not present at the same time.

Each analyst (subject) was given a series of simple "radar classification problems." Each problem involves two simple steps. Step 1, each analyst receives three parameter readings about an aircraft in the air space, such as: SPEED = Low, RANGE = Short, SIZE = Small. Step 2, on the basis of this information, each analyst classifies the aircraft as either: HOSTILE, NEUTRAL, or FRIENDLY. Information was given to the subjects and collected from them electronically. In Fig. 20.2 an illustrative display is shown. After the subject makes his or her decision, he or she is asked to provide an estimate of confidence in that decision. The subject's decision, speed, and confidence are stored.

The same procedure is repeated for the subjects acting as managers. There are two differences between the analyst and managerial condition. First, in the analyst condition the subjects receive raw information on the aircraft. Whereas, in the managerial condition the subjects receive the decisions of a set of nine human subjects. Second, in the analyst condition the subjects are told that their decision is not the final organizational decision, that they are working as part of a team, and that they are receiving reduced information about aircraft from scanners. In contrast, the subjects acting as

[4]The human experimental data reported on here is a subset of that collected by Carley and Prietula. The automated data collection procedure was written by Prietula.

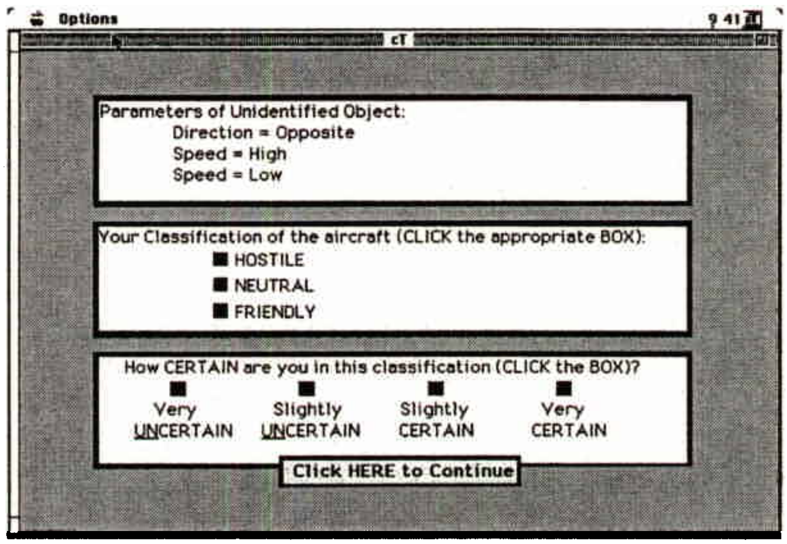

FIG. 20.2. Illustrative data collection screen.

managers are told that their decision is the final organizational decision, that they are receiving information from a set of nine analysts as to what those analysts think is the true state of the aircraft.

PERFORMANCE

Performance is measured at the organizational level. Each organization for each problem makes a single decision. For teams with voting the organizational decision is the majority vote of the group. For teams with manager the organizational decision is the decision made by the manager. For each problem there is a "correct" answer, the true state of the aircraft. Performance is measured in terms of the number of times the organizational decision is the correct decision.

Overall performance is defined as the percentage of decisions made by the organization that match the true state of the aircraft. For example, a match occurs if the organizational decision is that the aircraft is FRIENDLY and the true state of the aircraft is FRIENDLY. A second measure is "slight error"—the percentage of decisions that are one step different from the true state of the aircraft. This can happen, for example, when the true state of the aircraft is NEUTRAL, and the organization thinks it is either FRIENDLY or HOSTILE. A third measure is "severe error"—the percentage of the decisions that are two steps different from the true state of the aircraft. This can happen, for example, when the true state of the aircraft is FRIENDLY and the organization thinks it is HOSTILE. Table 20.1 shows the mapping between organizational decision, true state, and these measures. These three measures are calculated separately for each phase. Thus, the percentage in each phase is based on 30 decisions.

RESEARCH DESIGN

A series of organizations varying in their coordination structure and the cognitive structure of the agent were analyzed. A total of 12 organizations were examined, four coordination schemes (two organizational structures by two resource access structures) by three types of agents. All agents, regardless of the organization or cognitive abilities are pretrained on a random sample of 10 problems. Each organization was then faced with a series of 60 problems/aircrafts. These 60 cases are divided into two phases, each of which has 30 cases. The difference between the first phase and second phase is that during the first phase the agents are learning and during the second phase the agents are not learning. During the first phase the agents receive feedback and during the second phase they do not. The 60 problems are drawn randomly from the possible 19,683, with the constraint that for each set of 30 one third of them have a true state of FRIENDLY, another one third have a true state of NEUTRAL, and the final one third have a true state of HOSTILE. All organizations get the same set of problems. Analyses using ELM reveal that these problems are typical of the overall set of 19,683 problems; that is, performance is not significantly different on this set than it is for the overall set. This suggests that the results should generalize to others sets of problems for this task.

TABLE 20.1
Definition of Organizational Performance Measures

Organizational Decision	*True State*		
	FRIENDLY	NEUTRAL	HOSTILE
NEUTRAL	correct	one away	two away
NEUTRAL	one away	correct	one away
HOSTILE	two away	one away	correct

RESULTS

Overall the organizations examined make the correct decision 58.1% of the time. Further, these organizations make slight errors (off by one) 34.7% of the time, and a severe error (off by two) 7.2% of the time. The agent's nature does affect the organization's performance (see Table 20.2). First, all agents do better than simply guessing. Secondly, ELM agents are more similar to Human agents than are Soar agents. Further, cognitively sophisticated agents tend to make fewer severe errors. On the one hand, what these results are suggesting is simply that humans are not particularly well suited to this task, in general, and that information aids that admit keeping track of past performance (as is done in ELM or Soar) improve organizational performance. A supporting point is that ELM agents who were trained not just on the training set, but on all 19,683 possible tasks, have higher performance than any of the agents shown here. A corroborating point is made by researchers interested in technology to support group decision making (Olson, 1989). However, these information aids may aid in overall performance but they may serve to mask critical failures that humans appear to pick up on. On the other hand, and more to the point of this chapter, what these results are suggesting is that organizational performance can be dramatically affected by how the agents in the organization are modeled.

In Table 20.2 we see that organizations of Soar agents tend to outperform organizations of either ELM or Human agents. This is true even though there is no transference of learned productions among Soar agents. The reason has to do with the variance among the individual agents. Soar agents, because they are choosing among rules in a stochastic fashion, can see exactly the same information and yet make slightly different decisions. ELM agents, on the other hand, given similar experiences, will make exactly the same deci-

TABLE 20.2
Agents and Organizational Behavior

	Organizational Behavior		
Agent	*Overall Performance*	*Light Error (One Away)*	*Severe Error (Two Away)*
ELM	0.567 (0.497)	0.338 (0.474)	0.096 (0.295)
Soar	0.654 (0.477)	0.317 (0.466)	0.029 (0.169)
Human	0.521 (0.501)	0.388 (0.488)	0.092 (0.289)

sion given the same information. Human agents, see exactly the same information and yet can make very different decisions. Thus, these results are suggesting that slight variations in individual response given the same information can lead to improved organizational decisions; whereas, extreme variation in individual response can lead to worse organizational decisions.

Overall, these results suggest that researchers interested in exploring organizational behavior resulting from agent behavior need to be careful about how they model the agents. Such a conclusion is of little import, however, if the difference in organizational behavior is simply scaled by agent type. In other words, if regardless of the type of coordination scheme or of whether or not the agents are undergoing training, organizational performance for organizations of Soar agents was always proportionally higher than that for ELM agents, which was in turn always proportionally higher than that for Human agents, the impact of the agent model would be less critical to the study of organizational behavior. Whereas, if there are interaction effects among agent cognition, coordination, and training, then in fact the model of the agent is critical to the study of organizational behavior.

In fact, there are interaction effects. All organizations, regardless of the agents' cognitive capabilities exhibit poorer performance when there is a manager (see Fig. 20.3). This is due largely to the fact that as decisions move up the hierarchy information is lost. This effect is quite robust and has been discussed by numerous organizational theorists using terms such as *infor-*

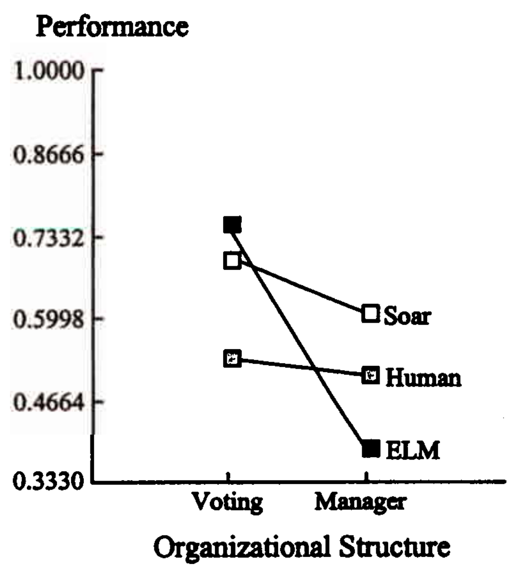

FIG. 20.3. Interaction between organizational structure and cognition.

mation condensation (Jablin, Putnam, Roberts, & Porter, 1986) and *uncertainty absorption* (March & Simon, 1958). Additionally, this analysis shows that organizations of ELM agents are relatively more disadvantaged by the manager than are Soar organizations which are relatively more disadvantaged than are human organizations. Part of ELM's relative disadvantage comes from the fact that the managers have rarely observed the patterns before and so are often guessing. However, organizations of ELM agents where the managers have seen all possible patterns still exhibit a greater reduction in performance when they have a manager than do organizations of other types of agents. The more cognitively sophisticated managers, Soar and human, do relatively better at this task than the ELM managers because they are not overcommitted to the lessons of history and can respond in a stochastic or "strategic" fashion. This ability, to overcome history through the fortuitous guess, causes organizations of these agents to be less disadvantaged by the information loss inherent in more hierarchical structures.

There are also interaction effects between the resource access structure and agent cognition. In particular, while the organizations of artificial agents tend to exhibit lower performance when information is distributed, organizations of humans tend to do better. Soar organizations are less disadvantaged by a distributed structure than are ELM organizations. The potential lack of consensus among analysts despite common incoming information results in worse decisions for Radar–Soar and ELM organizations; but, better decisions for human organizations. Further, Soar organizations are less

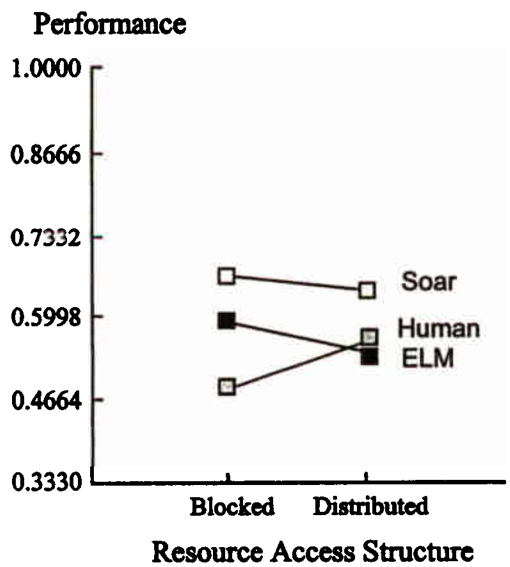

FIG. 20.4. Interaction between resource access structure and cognition.

disadvantaged by this lack of consensus, than are ELM organizations, as it occurs in both resource access structures. In the blocked structure the lack of consensus among Radar–Soar analysts is caused by interpretation (analysts seeing the same information but stochastically making different choices). In the distributed structure the lack of consensus among Radar–Soar analysts is caused by different information (analysts see different information and make different choices). For Soar agents the lack of consensus, coupled with the ability to respond stochastically results in overall better performance. For ELM, who learn slowly and respond consistently, the low level of training results in worse performance in a distributed environment because on average, more of the agents in the organizations will be guessing. In fact, had the ELM agents been fully trained (not shown) then the ELM organization, like the human organization, would have had an increase in performance when the resource access structure was distributed. Humans can take advantage of the greater resolution afforded by the distributed structure as can more fully trained ELM organizations.

Now consider these impacts only in organizations with managers. In Fig. 20.5, the relative impact of the different types of cognitive agents under these resource schemes for just teams with managers are shown. As note previously, ELM organizations when they have managers, the managers have less to learn in a blocked than in a distributed structure. As a result,

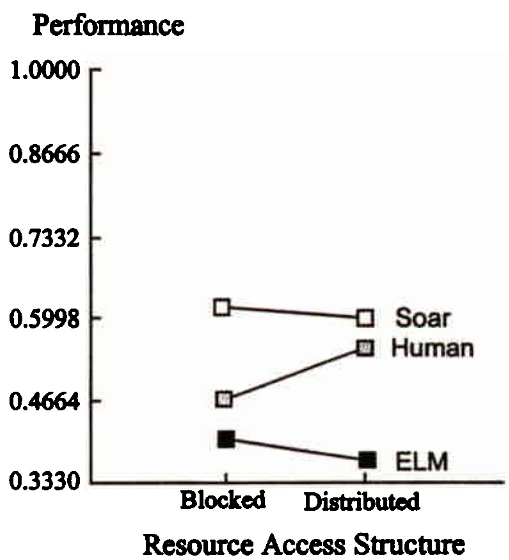

FIG. 20.5. Interaction between resource access structure and cognition in teams with managers.

they should and do, do better in this structure when they are only minimally trained as in this experiment. This lowering in performance for the ELM organizations is due simply to lack of training. Soar and human organizations can overcome this in two ways. First, differences in agent interpretation by the analysts (which should lead to the problem being "harder" from the manager's standpoint as it increases the number of patterns actually improves performance as the organization does not get trapped by an erroneous understanding of historical precedence as do the ELM organizations.

TRAINING

Now consider the impact of training. Both ELM and Soar organizations do better in Phase two when they are not receiving feedback than they did in Phase one when they did receive feedback. This is caused, not by some perverse ability to do better in the absence of feedback but simply because the only impact feedback to these artificial organizations into alter what decisions they have made. During phase one, the individuals in all organizations are learning. Feedback moves the agents, whether human, Soar, or ELM, out of the realm of guessing and improves their accuracy. Soar and ELM organizations do better in phase two as they remember the lessons of phase one and are less likely to guess during the second phase than the first. The

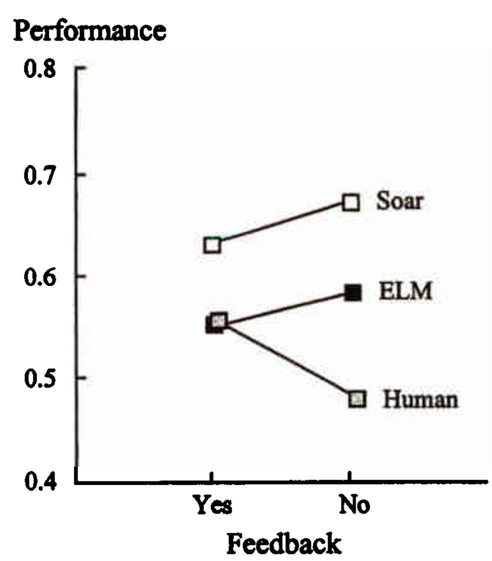

FIG. 20.6. Impact of training.

absence of feedback keeps them from learning more but does not inhibit them from applying the lessons of history. In contrast, within human organizations feedback appears to be necessary both for learning and for convincing the humans to continue to apply the lessons of history.

DISCUSSION

In interpreting the foregoing results the reader should keep in mind the following caveats. First, these results are based on only 30 problems per phase where the 30 problems were chosen at random from a set of 19,683 possible problems. This is a very small proportion of the possible problems. Future work should consider increasing the sample size, or running the overall study multiple times with different sample problems.

Second, the results for each type of agent are based on a single organization. For ELM, this is not a problem as all ELM organizations faced with the same set of problems in the same order will respond in the same way. For Soar, this is more of a problem as the stochastic nature of the decision making will lead to slight differences in performance across different Soar organizations. For Humans, this is yet more of a problem. That is, we expect that there is even greater variance in the behavior of individual humans than there is in behavior of Soar agents and certainly more than in the behavior of ELM agents. Future studies should run a Monte-Carlo analysis for Soar organizations and should average the behavior for multiple Human organizations.

Third, the results presented have been at the organizational level. They provide no insight into whether at the individual level Soar agents acts more like Human agents than do ELM agents. Future work should explore this issue through a more detailed analysis of individual action. Such an exploration could look at both individual performance, and also at the impact of the order of decisions on the specific learning pattern achieved by the various types of agents. Such analyses would provide greater insight into how agent level actions change in response to cognition and coordination.

Fourth, the radar task that we used is highly stylized and differs from a real radar task in interesting ways. For example, there is no autocorrelation among aircraft features. Whereas, in reality these features tend to clump. For example, aircraft with weapons emission signals also are often traveling very fast. At issue is whether this lack of clumpiness makes the problem relatively more difficult for humans than for the simulated agents. The data we have collected cannot completely address this question. However, when the presence of features is not uniform (as in this study) and when sets of features tend to go together in predicting some outcome the problem is "easier" for humans. Certainly, for ELM agents, such clumpiness or bias reduces uncertainty and increases performance (Carley & Lin, 1995). Given the Soar agent model, it will also be the case that for Soar agents clumpiness will improve performance. Because the performance of all agents will improve in such a situation, it is not clear whether in fact humans will be relatively more advantaged. Future research might address this question.

All these caveats aside, the import of this study is that it demonstrates that the model of cognition affects organizational performance results in complex and interesting ways, and that interactions between the organizational structure and the model of cognition are important in determining overall organizational performance. This research on coordination using computational models differs from other work in this venue in two ways. First, studies of coordination that employ computational models of agents rarely contrast the behavior of the simulated agents with humans (Durfee, 1988). Such a contrast is important when the goal is predicting actual organization behavior. Herein, a basic attempt at making such comparisons is made. Second, many of the computational studies of coordination employ large models that emulate the organizations in question (Gasser & Majchrzak, 1992, 1994; Levitt et al., 1994) rather than examine the behavior of highly stylized coordination schemes such as those examined here. The models examined herein are less "accurate" in describing organizational coordination, but are better models of individual cognition. Important advances in understanding organizational performance can be made by contrasting the results from the two styles of computational modeling. Future studies should consider extending these results by contrasting alternate models of individual behavior and by attending to the limits of the current study.

CONCLUSION

This research suggests that complex adaptive agent models can be fruitfully applied to the study of organizations. Two different complex adaptive agent models of humans were constructed, these agent models were placed within models of organizational coordination, and the resulting organizational performance was examined. These simulated results were compared with the actual performance of humans engaged in the same coordination structures. This combination of simulation and human experiments is particularly valuable in understanding the veridicality of the computational models. Moreover, this combination is valuable in understanding how the veridicality of the agent model interacts with structural and task changes in affecting organizational performance. Consider the following two examples.

Clearly, the forgoing analysis supports previous research that teams outperform hierarchies (regardless of the type of agent) and that organizations of artificial agents often outperform organizations of humans. More importantly, this analysis refines these points by pointing to important interactions between structure and cognition. For example, for managers, the ability to respond stochastically, to make "strategic" decisions, improves performance. The result of these "guesses," is that differences in performance due to the organizational structure are minimized when the agents in the organizations are more cognitively capable. The fortuitous guess has advantages. This compliments earlier work by Carley and Lin (1995) showing that factors facilitating managerial decision making, such as bias in the task, serve not only to improve performance but to mitigate the impact of structure. Ouchi (1980) suggested that as task complexity and information uncertainty increases new forms of coordination may be needed. This chapter, in contrast, suggests that new coordination qua mechanisms for linking individuals together, may have little impact as the agents in these organizations are generally intelligent. Performance should become not just a function of the coordination scheme linking individuals (Malone, 1986); but also, the specific messages being communicated among linked positions. Thus improvements in coordination might be achievable not through structure, but through content.

Previous research suggested that specialization, much as we see in the blocked structure, is disadvantageous. For instance, in blocked structures agents develop unique skills, perspectives, and ways of thinking (Brewer & Kramer, 1985) that can reduce the group's flexibility in a crisis (Carley, 1991) and exacerbate conflict when it occurs (Dearborn & Simon, 1958; Jablin, 1979; Monge, Rothman, Eisenberg, Miller, & Kirste, 1985). This suggests that consensus is valuable to the organization. In this chapter, however, it was found that for organizations the lack of consensus among agents caused by interpretation (different decision on same information) results in improved performance; but the lack of consensus among agents caused by resolution (seeing different information) tends to degrade performance. In other words a shared perspective, whether or not it results in consensus seems to be important in increasing performance. However, whether the observed results are actually the result of having a shared perspective, or a function of the degree of disagreement, needs further study.

Complex interactions among agent cognitive capacity, task complexity, and organizational design affect overall organizational performance. The implication of this finding is that researchers must be very careful in interpreting results of studies of organizational performance based on modeling organizations as collections of intelligent adaptive agents as such results maybe highly dependent on the agent model.

ACKNOWLEDGMENT

This work was supported in part by Grant No. IRI-9111804 by the National Science Foundation.

REFERENCES

Brewer, M. B. and R. M. Kramer (1985). The psychology of intergroup attitudes and behavior, *Annual Review of Psychology, 36*, 219–243.

Bush, R. R. and F. Mosteller (1955). *Stochastic Models of Learning*, New York: Wiley.

Carley, K. M. (1990). Coordinating for Success: Trading Information Redundancy for Task Simplicity. *Proceedings of the 23rd Annual Hawaii International Conference on System Sciences*.

Carley, K. M. (1991). Designing Organizational Structures to Cope with Communication Breakdowns. *Industrial Crisis Quarterly, 5*(1), 19–57.

Carley, K. M. (1992). Organizational Learning and Personnel Turnover. *Organization Science, 3*(1), 2–46.

Carley, K. M. and J. Harrald (1993). Organizational Learning Under Fire: Theory and Practice. *American Behavioral Scientist, 40*(3), 310–332.

Carley, K., J. Kjaer-Hansen, M. Prietula, and A. Newell (1992). Plural-Soar: A Prolegomenon to Artificial Agents and Organizational Behavior. In M. Masuch & M. Warglien (Eds.), *Artificial Intelligence in Organization and Management Theory* (pp. 87–118). Amsterdam, The Netherlands: Elsevier.

Carley, K. M. and Z. Lin (1995). Organizational Designs Suited to High Performance Under Stress. *IEEE Systems, Man and Cybernetics, 25*(1).

Carley, K. M. and Z. Lin (1997). A Theoretical Study of Organizational Performance under Information Distortion. *Management Science, 43*(7), 976–997.

Carley, K. M. and A. Newell (1994). The Nature of the Social Agent. *Journal of Mathematical Sociology, 19*(4), 221–262.

Carley, K. M., D. Park, and M. Prietula (1993). Agent Honesty, Cooperation and Benevolence in an Artificial Organization. *Proceedings of AI and Theories of Groups and Organizations: Conceptual and Empirical Research*, Eleventh National Conference on Artificial Intelligence, Washington, DC.

Carley, K. M. and M. J. Prietula (1994). ACTS Theory: Extending the Model of Bounded Rationality, In K. M. Carley and M. J. Prietula (Eds.), *Computational Organization Theory* (pp. 55–87). Hillsdale, NJ: Lawrence Erlbaum Associates.

Carroll, G. R. (1984). Organizational Ecology. *Annual Review of Sociology, 10*, 71–93.

Cohen, M. D., J. B. March, and J. P. Olsen (1972). A Garbage Can Model of Organizational Choice. *Administrative Science Quarterly, 17*(1), 1–25.

Dearborn, D. C. and H. A. Simon (1958). Selection Perception: A Note on the Departmental Identification of Executives. *Sociometry, 21*, 140–144.

Durfee, E. H. (1988). *Coordination of Distributed Problem Solvers*. Boston, MA: Kluwer.

Galbraith, J. (1973). *Designing Complex Organizations*. Reading, MA: Addison-Wesley.

Galbraith J. R. (1977). *Organization Design*. Reading, MA: Addison-Wesley.

Gasser, L. and A. Majchrzak (1992). HITOP-A: Coordination, Infrastructure, and Enterprise Integration. *Proceedings of the First International Conference on Enterprise Integration*. Boston: MIT Press.

Gasser, L. and A. Majchrzak (1994). ACTION Integrates Manufacturing Strategy, Design, and Planning. In P. Kidd and W. Karwowski (Eds.), *Ergonomics of Hybrid Automated Systems IV*. Netherlands: IOS Press.

Hollenbeck, J. R., D. R. Ilgen, D. J. Sego, J. Hedlund, D. A. Major, and J. Phillips (1995). The Multilevel Theory of Team Decision Making: Decision Performance in Teams Incorporating Distributed Expertise. *Journal of Applied Psychology, 80*, 292–316.

Jablin, F. M. (1979). Superior-subordinate communication: The state of the art. *Psychological Bulletin, 86*, 1201–1222.

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

Laird, J. E., P. S. Rosenbloom, and A. Newell (1986). *Universal Subgoaling and Chunking: The Automatic Generation and Learning of Goal Hierarchies*. Boston, MA: Kluwer.

Laird, J., A. Newell, and P. Rosenbloom (1987). Soar: An Architecture for General Intelligence. *Artificial Intelligence, 33*, 1–64.

La Porte, T. R., and P. M. Consolini (1991). Working in Practice But Not in Theory: Theoretical Challengers of 'High-Reliability Organizations'. *Journal of Public Administrative Research & Theory, 1*(1), 19–47.

Levitt, R. E., G. P. Cohen, J. C. Kunz, C. I. Nass, T. Christiansen, and Y. Jin (1994). A Theoretical Evaluation of Measures of Organizational Design: Interrelationship & Performance Predictability. In K. M. Carley & M. J. Prietula (Eds.), *Computational Organization Theory* (pp. 1–18). Hillsdale, NJ: Lawrence Erlbaum Associates.

Lin, Z. and K. M. Carley (1993). Proactive or Reactive: An Analysis of the Effect of Agent Style on Organizational Performance. *International Journal of Intelligent Systems in Accounting, Finance and Management, 2*, 271–288.

Lin, Z. and K. M. Carley (1997). Organizational Response: The Cost Performance Tradeoff. *Management Science, 43*(2), 217–234.

Mackenzi, K. D. (1978). *Organizational Structures*. Arlington Heights, IL: AHM Publishing Corporation.

Malone, T. W. (1986). Modeling Coordination in Organization and Markets. *Management Science, 33*(10), 1317–1332.

Masuch, M. and P. LaPotin (1989). Beyond Garbage Cans: An AI Model of Organizational Choice. *Administrative Science Quarterly, 34*, 38–67.

Mintzberg, H. (1979). *The Structure of Organizations*. Englewood Cliffs, NJ: Prentice-Hall.

Mitchell, W. M. (1988a). Toward a Unified Theory of Cognition. *Science, 241*, 4861, 27–29.

Mitchell, W. M. (1988b). Soar: A Unified Theory of Cognition? *Science, 241*, 4863, 296–298.

Monge, P. R., L. W. Rothman, E. M. Eisenberg, K. L. Miller, and K. K. Kirste (1985). The dynamics of organizational proximity. *Management Science, 31*, 1129–1141.

Newell, A., G. Yost, J. E. Laird, P. S. Rosenbloom, and E. Altmann (1991). Formulating the Problem-Space Computational Model. *25th Anniversary Commemorative of Computer Science*, Carnegie Mellon University.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Olson, M. (Ed.). (1989). *Technological Support for Work Group Collaboration*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Ouchi, W. G. (1980). Markets, bureaucracies, and clans. *Administrative Science Quarterly, 25*, 129–140.

Papageorgiou, C. P. and K. M. Carley (1992). A Cognitive Model of Decision Making: Chunking and the Radar Detection Task. *CMU-CS Technical Report*.

Roberts, K. (1989). New Challenges to Organizational Research: High Reliability Organizations. *Industrial Crisis Quarterly, 3*(3), 111–125.

Roberts, K. (1990). Some Characteristics of One Type of High Reliability Organizations. *Organization Science, 1*(2), 160–176.

Rosenbloom, P. S., J. E. Laird, A. Newell, and R. McCarl (1989). A Preliminary Analysis of the Soar Architecture as a Basis for General Intelligence. *Proceedings of the workshop on Foundation of Artificial Intelligence*, Cambridge, MA: MIT Press.

Thompson, J. (1967). *Organizations in Action*. New York: McGraw Hill.

Williamson, O. E. (1975). *Market and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.

Ye, M. and K. M. Carley (1995). Radar-Soar: Towards An Artificial Organization Composed of Intelligent Agents. *Journal of Mathematical Sociology, 20*(2–3), 219–246.