Annals of Operations Research 75(1997)25–47

Organizational adaptation*

Kathleen M. Carley

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

E-mail: Kathleen.Carley@centro.soar.cs.cmu.edu

A computational model of organizational adaptation in which change occurs at both the strategic and the operational level is presented. In this model, simulated annealing is used to alter the organization's structure even as the agents within the organization learn. Using this model a virtual experiment is run to generate hypotheses which can be tested in multiple venues. The results suggest that, although it may not be possible for organizations of complex adaptive agents to locate the optimal form, they can improve their performance by altering their structure. Moreover, organizations that most successfully adapt over time come to be larger, less dense, with fewer isolated agents, and fewer overlooked decision factors. These results have implications for organizations of both humans and non-humans. For example, they suggest that organizational learning resides not just in the minds of the personnel within the organization, but in the connections among personnel, and among personnel and tasks. These results suggest that collections of non-humans may come to seem more intelligent (i.e., show improved performance) even if the agents remain unchanged if the system simply develops duplicate copies of some of the artificial agents and if the connections among agents are dynamically altered.

1. Introduction

Literature on organizational adaptation suggests that organizations change over time [19,49,55]. Part of this change is due to strategic re-organization [7,32], including re-engineering and re-organization. However, not all types of re-organization are equally valuable. For example, organizational performance may improve as individual members of the organization gain experience [41]. But, organizational downsizing may lead to corporate anorexia as the organization eliminates personnel and so loses the benefits of their experience. Indeed, organizational theorists are faced with many

* This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-93-1-0793 (UConn FRS 521676). The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. government.

questions concerning adaptation. For example, what leads to successful adaptation? Or, more specifically, do organizations that adapt successfully have different organizational designs? Do such organizations follow different patterns of adaptation?

Little is known about organizational adaptation, and even less about how organizations should change. Most theories of organizational design speak to the relative advantage of different designs in different situations [24,36]. Such theories, in principle, provide some guidance for organizational change. For example, population ecology can be interpreted as suggesting that if the organization is moving out of a niche environment, then the organization should move from a specialist to a more generalist structure [24]. As another example, Staw et al. [54] have argued that organizations, when faced with a decrease in their performance, will shift to a more rigid and centralized structure such as is typical in many hierarchical forms. Such suggestions, however, provide little theoretical guidance as to the path of change, or the relative advantages and disadvantages of different adaptation strategies.

The dynamics of change result from simple, but possibly nonlinear processes. Consequently, thinking through the implications of adaptation processes is non-trivial. Consider the following two illustrative processes which may occur simultaneously. When performance drops, organizations may enter a downward spiral by choosing to lay personnel off, thereby losing experience, which in turn may lead to a further reduction in performance, which may lead to further downsizing. Alternatively, such layoffs may lead to a reduction in non-essential personnel, thereby freeing managerial time to attend to the decisions of key personnel, thus making the task simpler for these managers, allowing them to learn faster, and increasing overall performance. Given just these two processes, what will be the impact of downsizing? How can issues such as these be addressed?

Recent advances in computational analysis and distributed artificial intelligence (DAI) suggest a possible avenue for theory creation in this complex domain of organizational adaptation. Researchers have begun to use complex adaptive agent models, such as genetic algorithms and neural networks, to answer questions about change. However, such analyses have typically focused more on the evolution of industries and the sets of organizations within a market, rather than adaptation within a single organization [2,4,18,25,26,46]. Another line of research, also employing computational models to explore organizational performance, is the work on organizational learning. These models typically examine either individual level [14,39,58] or organizational level learning [35], rather than examining the interaction among the two levels of learning [11,33]. These two streams of research are complemented by the work within DAI. In this case, researchers have examined organizational adaptation by focusing on the effect on performance of coordination and communication among intelligent agents [20,31,37], planning [23], monitoring [22], and socially shared cognition [29,30]. Often this research focuses on the decision making capabilities of only a few agents and there is little attention to organizational design and factors such as the size of the organization, the number of levels in the hierarchy, and so on.

Analyses, such as those mentioned above, have demonstrated the power of computational models for theory building and for examining issues of social and organizational dynamics. Such models abstract away many of the complications existent in actual organizations and so lay bare the relationships among the remaining components of organizational design and adaptation. Further, these computational models can, and generally do, contain complex nonlinear processes. In fact, computational analysis is one of the few techniques that enables the theorist to think through the possible ramifications of such processes and to develop a series of consistent predictions. Consequently, computational models can be, and have been, used in a normative (and sometimes a prescriptive) fashion to generate hypotheses that can then be tested in other empirical settings. Researchers taking this approach use these models to run virtual experiments [1] and so generate a series of hypotheses which can then be tested in other empirical settings.

From the perspective of a science of organizing, these computational models speak to another point. That is, these computational models provide basic information about organizing. Indeed, researchers have long used the human organization as a metaphor for the organization of computational (and even cognitive) processes (e.g. [44]) and they have treated multi-agent models as computational analogs of human organizations [15]. Computational models allow researchers to show proof of concept and to demonstrate that certain factors, which can be completely modeled, can or cannot generate certain phenomena. These computational models employ the use of "artificial" agents, acting as humans. Thus, the predictions these models make may be applicable to organizations of humans, and perhaps may be equally applicable (and some would argue more so) to organizations of "non-humans". Depending on the way in which the agents are modeled, the results of these models may be interpreted as predictions about organizing in general. In this way, multi-agent computational models are a theory building tool for researchers interested in organizations broadly speaking (whether composed of humans, non-humans, or collections of the two) and in the process of organizing.

In this paper, a computational approach is taken to the issue of organizational adaptation. Drawing on various literature previously discussed, a dual-level model of

[1] A virtual experiment is an experiment where the "agents" are simulations, not human beings or animals. The term simulation has been used to refer both to the program (the computational model) and the result of running that computational model at least once. With modern computational models, the space of options is sufficiently large that rarely can the complete response surface of the model be calculated. Thus, one approach to locating the model's predictions is to run a virtual experiment; that is, to vary a small set of parameters within a specific experimental design and then analyze the results statistically. A virtual experiment is a computational analog of a laboratory experiment, subject to similar issues in design and statistical analysis of the results. Virtual experiments can be, and are being, used to aid the researcher in pre-testing the design of human experiments (by running the experiment with computer models that instantiate the theory instead of humans) before they are run in a lab. Virtual experiments also play a critical role in the development of computational theories as they admit hypothesis generation.

organizational adaptation is presented in which the organization can change at both the strategic and the operational level. At the operational level, the organization is modeled as a collection of adaptive agents, each of whom occupies a particular organizational position and has the capability of learning over time as they gain experience with the task they are performing. Agents are modeled essentially as a Bush and Mosteller [6] stochastic learning model with additional limits on attention, memory, and information processing which effectively bound the agent's rationality far beyond those in the original stochastic models. At the strategic level, the organization can adapt strategically in response to changes in its performance by altering its design in a number of different ways including downsizing, expansion, and re-engineering. This strategic adaptation is modeled as a simulated annealing process. Using this computational model, a series of virtual experiments will be done to address the question "what leads to successful adaptation?" This model has been informed by empirical studies both on individual learning by humans and on adaptation within human organizations. Nonetheless, since this model portrays the agents as abstract entities capable of doing only one task and learning only in a limited fashion, the results can be thought to apply equally to organizations of humans and non-humans.

In presenting this model, a somewhat agnostic stance is taken with respect to whether this is a model of human organizations or of non-human organizations. The model is, simply, a model of organization; i.e., a system that can adapt and that is filled with adaptive agents. The results from a virtual experiment that focuses on the impact of organizational design on performance and the ability of the organization to successfully adapt are presented. These results are then interpreted as more specific hypotheses for both organizations of humans and non-humans. However, before presenting the model and discussing some of its implications, the basic rationale for modeling organizational strategic adaptation as a simulated annealing process is presented.

2. Organizations as simulated annealers

Why might organizational theorists be interested in using it as a model of organizational strategic adaptation? The basic argument is quite simple. Simulated annealing can be interpreted as a computational analog of the imperfect optimization process organizations appear to go through when they alter their design in an attempt to improve performance. In a detailed empirical study of investment banking, Eccles and Crane found that the process of strategic change in organizational design gone through by human organizations appears to be an annealing process [21]. Perhaps of equal importance to a science of organizing, the process of optimizing the organizational design for an organization of agents is so complex a task that a heuristic approach, like simulated annealing, is called for. But what exactly is simulated annealing?

Simulated annealing, a heuristic approach to optimization, is intended to be a computational analog of the physical process of annealing a solid (see [34], and for an overview, see Rutenbar [50]). The goal of the annealing process is to find that state (atomic configuration) which minimizes system costs (energy). The process of annealing involves heating the system to a state that admits many alterations, then, given a schedule of decreasing temperatures, cooling the system slowly so that it reaches thermodynamic equilibrium at each temperature in this schedule, and eventually freezing the system in a good configuration. This process is carried out by having a set of possible moves for altering the existent state to another state, choosing a move, evaluating the proposed state that this move would create, and then moving to that new state if it improves things and possibly even if it does not. Further, the frequency with which such non-improving moves are accepted decreases with time (as the temperature cools).

Pictorially one might imagine a three-dimensional surface in which all states are arrayed on two dimensions with the height of the surface being the cost of that particular state. The cost function plays a critical role as it determines the shape of the surface. Simulated annealing has been likened to a process of using a ball to locate the lowest point on such a surface. Imagine randomly dropping a single ball on this surface, rolling it around, and then shaking the surface to pop the ball to a new location (with decreasing frequency over time). The spot on the surface where the ball touches is the current state of the system. Like the ball on this surface, the system can be in only one state at a time. Clearly, some surfaces will be easier for the ball to traverse and find the lowest spot. For example, single peak surfaces are relatively easy to traverse.

Heuristic optimization techniques such as simulated annealing are not guaranteed to find the optimal solution; however, they do satisfice. That is, they move the system to a state that is typically better than the initial state but may not be the best. Returning to the pictorial description, if the surface described by the cost function is extremely lumpy, then there is no guarantee that shaking the surface will move the ball to the lowest point; i.e., there is no guarantee that the annealer will locate the lowest cost state. Further, for combinatorial optimization problems which are NP-complete it may not be possible to locate the exact solution in a reasonable amount of time. Thus, heuristic solutions like simulated annealing are often the only practical answer.

However, how does this apply to organizations? On the one hand, simulated annealing is a reasonable approach to the organizational design problem. Organizations can be viewed as faced with the design problem; i.e., the need to locate the organizational design that optimizes organizational performance subject to various constraints. Organizational performance is a function of a large number of factors of which the various elements of design is only one component, but one over which the organization has some, albeit limited, control. Thus, the organizational design problem is, at least, NP-complete. From a purely technical perspective, some type of heuristic-based approach appears to be called for.

On the other hand, simulated annealing is a reasonable computational analog of organizational strategic adaptation. Let us consider this argument in greater detail. To begin with, there is a significant difference between organizations and general collectivities of agents, to wit, within the organization there is a CEO or central unit that directs some of the change in the way the set of agents is structured or connected, which agents are connected, and which agents are doing what. These connections and assignments can be thought of as the organization's design. The organization moves through a series of organizational designs, one at a time. This directed change in the organization's design can be characterized as strategic adaptation. The organization (specifically, the CEO or central unit) has a set of possible strategies (move set) that dictate which designs are possible given the current design. Such strategic adaptation requires the strategist (the CEO or central unit) to have knowledge about which agent knows what, which agent has which capabilities, and it requires the ability to anticipate, however imperfectly, the future.

Over time, the organization attempts to optimize its design given some cost function. For organizations, the cost function depends on what the organization perceives as important, e.g., minimizing salary, maximizing the number of widgets produced, or maximizing decision accuracy. Further, the organization's optimization process is imperfect. That is, a strategic change is employed if it appears to move the organization closer to the goal regardless of whether or not it actually does so and regardless of whether or not it is the best change for accomplishing that goal [42,51]. The organization (more precisely, the CEO or central unit) is not omniscient, does not compare all strategies, but simply evaluates a strategy through a kind of "what if" analysis, trying to forecast or anticipate, albeit imperfectly, the future [1,3,16]. Since the forecast is known to be imperfect, the organization may at times gamble on re-designs that might possibly "increase costs" if it is felt that there is some long-term advantage. In general, organizations are much more prone to this kind of risky behavior when they are new. This can be characterized as the liability of newness [56]. As the organization matures, the number of high-risk moves decrease, and the organization becomes more staid and locked into a certain way of doing business. This staidness has also been characterized as competency traps [38].

The parameters and processes of simulated annealing have direct and obvious translations into known organizational behaviors. In this sense, an annealing model of organizational change has reasonable face validity. The mapping of simulated annealing onto organizational strategic adaptation is summarized in table 1. Even at this gross level, there is a reasonable mapping between simulated annealing and organizational strategic adaptation. In actuality, the mapping exists even at a finer level of detail. Some of this additional detail will be seen when the computational model is described.

3. Computational model of organizational adaptation

A dual-level information processing model of an organization adapting in response to its environment is employed (see figure 1). The organization acts at both

Table 1

Comparison of simulated annealing and organizational adaptation.

Simulation annealing	Organization strategic adaptation
System	Organization's CEO or central unit
State	Organizational design
Current state	Current organizational design
Temperature	Risk aversion
Accepting a cost increasing move	Taking a risk
High temperature means accepting many cost cost increasing moves	Liability of newness
Move set	Re-design strategies
Heuristic optimization process	Imperfect optimization process
Minimize cost	Maximize performance
Cooling schedule	Approach to becoming risk averse
Proposed state	Proposed new design
Evaluation of proposed state	Limited lookahead, anticipation of future
State evaluation	Observed performance

a strategic and an operational level. At the operational level, organizational performance is determined by the actions of the individuals in the organization as they work on tasks. The specific model used is the CORP model of organizational performance [12, 14]. At the strategic level, organizational performance is affected by the ability of the CEO or central unit to anticipate the future and take the appropriate strategic actions to alter the organization in response to environmental cues. The model of strategic adaptation is based on a simulated annealing model.

3.1. Operational level: CORP

CORP is a simple information processing model of organizational performance in which organizational learning is the aggregate of individual learning plus learning the appropriate weights (or trust) in other's decisions [14]. Within CORP, organizational performance is seen as a function of the type of training received by the agents (experiential or procedural), the type of task (complexity – number of choices, bias – likelihood of certain outcomes, and decomposability – interdependence among subparts), the reporting structure (who reports to whom), S, and the resource access structure (who has access to what resources or information), R. CORP, related models, and predictions from the CORP model have been extensively described in previous studies [9, 12, 14]. CORP has been shown to be a reasonable model of organizational performance, both against experimental lab studies [13] and archival data on actual organizations [9, 40]. Models like CORP, or extensions of CORP, have received

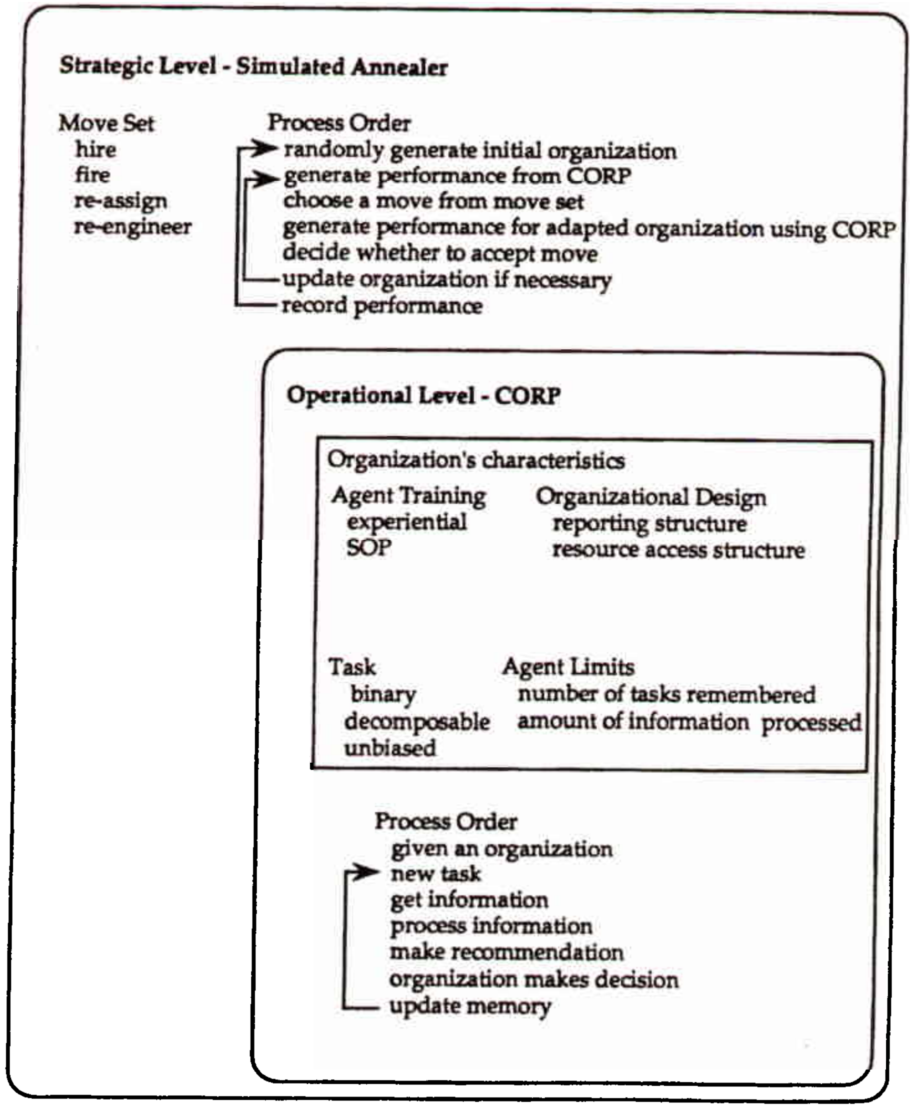

Figure 1. Overview of dual-level model.

extensive attention [43,45], as has the binary and trinary choice tasks underlying CORP (see for example, [12,14,27,28,47,57]).

In CORP, there is a single organization composed of a set of intelligent adaptive agents, each of whom must process some task-related information, and all of whom are connected within a single reporting structure. The performance of this single organization is its accuracy given a classification task, which is measured as the percentage of problems in a window of opportunity that the organization correctly classifies. Basically, the organization is faced with a sequence of tasks drawn from at random with replacement from the set of possible tasks. The specific task used for this paper is a nine-bit binary unbiased decomposable classification task in which the true decision is an A if there are more 1's than 0's in the set of task features and B otherwise. Each time period, the organization must make a decision for the current task. The organizational decision making process involves a series of individual decision-making processes. Of the $N(t)$ agents who are collectively processing $I(t)$ resources or pieces of information, some are processing the information on that task and are reporting to other agents, while other agents are simply processing the reports of others.

The agents in the organization are all either experiential learners [14] or they follow standard operating procedures [12]. Each experiential agent classifies the pattern of information that it sees and recommends the choice that was most often correct in the past for that pattern. If the agent has no previous experience, it simply guesses. This information can be either or both raw information on this specific task or the recommendations of other agents for this task. After the organization makes its decision, each agent receives feedback as to what was the correct choice for that task. Each agent then increments its memory. Each procedural agent simply follows this standard operating procedure: report choice A if there are more 1's than 0's in the pattern it observes, and choice B otherwise. With sufficient experience, experiential agents typically come to resemble procedural agents. Changes in the organizational design may alter the number of pieces of information various agents see for both experiential and procedural agents.

In keeping with the information processing perspective, all agents are boundedly rational both in terms of organizational access to information and in terms of cognitive ability to process information [8,52]. First, agents can only handle a maximum of seven pieces of information. Second, agents do not remember exactly what happened on a particular task; rather, they remember general trends. Third, what information the agents know is a function of their position in the organization. Fourth, experiential agents have an additional limitation which is that they do not remember all the tasks that they have seen. Rather, they suffer from both a primacy and a recency effect and thus remember only the first 500 tasks and the most recent 500 tasks they have seen.

The organization may be composed of anywhere between 2 and 45 agents organized in one to three layers, with a maximum of 15 agents per layer. The organizational decision is the CEO's decision. The CEO's decision is based on the recommendations of the agent or agents in the top tier in the organization. Essentially, the CEO simply makes that decision recommended by the majority of the top tier managers. If there is no majority decision, then the CEO randomly chooses one of the two options.

3.2. *Strategic level: Simulated annealing*

Organizational strategic adaptation is modeled as a simulated annealing process, such that the organization's restructuring strategies are the move set. The move set includes the following actions: *firing* – drop n agents (such that $1 \leq n \leq N_0(t)$), *hiring* – add n agents (such that $1 \leq n \leq N_{max} - N_0(t)$), *re-assigning* – delete the tie between agent i and j (i reports to j) and reassign agent i to report to agent k, and *re-engineering* – delete the tie between agent i and resource s and add a tie between agent j and piece of information s. Exactly how many agents are hired (or fired, re-engineered, or re-assigned) at a time is given by a Poisson distribution.[2] Notice that the effect of the connection changes (re-assignment and re-engineering) is to simply move connections and will not lead to an absolute increase or decrease. This type of connection-based change was used so as to distinguish simple tie movement from the more extensive tie changes caused by adding or dropping nodes.

The organization begins with a particular design ($S(0)$ and $R(0)$) and proceeds to process 500 tasks. After this, its basic life cycle begins. First, the performance of the organization for a sequence of 100 tasks is generated using CORP, then a move from the move set is chosen and a new organizational design is suggested. This design is then hypothetically evaluated in a limited lookahead for 100 tasks, then the forecasted performance of the proposed design is compared with the previous performance of the current organization and a strategic decision is made as to whether or not to accept the change. Finally, if the change is accepted the organization's design is altered and the process begins again, whereas if the change is not accepted, the process begins again with the unchanged organization. Performance at time t for the current organization is the percentage of most recent 500 tasks that the organization correctly classified prior to time t.

The probability of accepting the new design is determined via the Metropolis criteria. Specifically, the change is always accepted if the forecasted performance for the hypothetical organization is better than the known performance of the current organization. Further, when the forecast is poorer, the change may still be accepted. In fact, we can think of the probability of accepting the "bad" design as the organization's degree of risk aversion. This probability is calculated, using the Boltzmann equation, as $P = P_0 e^{(-\Delta \operatorname{cost}(t)/T)}$, such that $cost(t) = 0 - \text{performance}(t)$ and P_0 is the initial probability of accepting a "bad" design. This probability decreases as the temperature decreases. Temperature drops every 100 tasks (time periods) as $T(t+1) = a * T(t)$, where a is the rate at which the organization becomes risk averse.

[2] Different strategic approaches than that explored herein can be characterized either by adding additional moves or by altering the mean value for the Poisson distributions for each of these four types of moves. One might also think of altering the means as simply a perturbation on a single strategy. A reasonable approach would be to consider small differences in the means (those that are not significantly different) as a procedure for representing perturbations of the same strategy, and large differences in the means (such as setting the mean to 0 versus the mean to 1) as a procedure for representing different strategies.

4. Virtual experiment

In order to examine the relationship between design and organizational adaptation, the following virtual experiment was run. Organizations were composed of either all experientially or all procedurally trained agents. Three types of strategic adaptation were considered – agent change (individual agents could be hired or fired), linkage change (individuals could be reassigned to new managers or tasks could be re-engineered and so components of the task are assigned to different agents), and general change (both agents and linkages). The size of the organization (2 to 45 agents), the number of levels (1 to 3), the initial reporting structure ($S(0)$), and the initial resource access structure ($R(0)$), were all chosen randomly with replacement from the set of possibilities. Each organization was simulated for 2000 time periods (after the initial 500). A total of 1000 organizations were sampled for each of the six condition (two types of training by three types of adaptation).

In exploring the impact of organizational structure on performance, a number of factors will be considered. The specific factors that will be attended to are: size, density, number of isolates, redundancy, and number of factors overlooked. Size, the number of agents in the organization is considered as, in principle, as the size of the organization is increased, more aspects of the task can be analyzed but the coordination problems increase. Density, the fraction of possible connections in the reporting structure that actually exist, is considered as the higher the density the higher the management workload but the greater the communication and so potential for noticing errors. The number of isolates is the number of agents in the organization that are not reporting to any other agent. Isolates represent points of organizational inefficiency as the work of these agents does not contribute to overall organizational performance. Redundancy, the number of agents per task for only those tasks that at least one agent is attending to, is examined as organizations are thought to need redundancy for error checking in complex tasks. This measure is calculated as the average number of agents working on a task factor and includes only those factors that are not being overlooked. Finally, the number of factors overlooked, like the number of isolates, represents a type of organizational inefficiency. A factor is overlooked if no agent is examining that factor. The higher the number of overlooks, the lower the organization's performance should be as information needed to make that decision is not being factored into the organization's decision. Finally, organizational performance is simply the fraction of tasks in the last 500 tasks seen by the organization that it correctly classifies. Herein, performance is often recoded into a six-point scale. Let x be the fraction of tasks that the organization correctly characterizes. Then the scale is as follows: less than or equal to 70% (= 0), $70\% < x \leq 75\%$ (= 1), $75\% < x \leq 80\%$ (= 2), $80\% < x \leq 85\%$ (= 3), $85\% < x \leq 90\%$ (= 4), and $90\% < x \leq 100\%$ (= 5).

For the binary choice task, the particular levels of training (500 tasks) and experiential memory (500 tasks) used tends to result in all agents, on average, being completely trained. For the binary choice task examined, in 500 trials each agent, even if it attends the maximum possible pieces of information (7) will on average have seen each pattern 4 or more times, which is sufficient for the agent to, on average, act like a majority classifier. When this is coupled with the fact that most agents will remain in the organization longer than 100 time periods and the fact that individual agents can remember up to 1000 pieces of information (500 through training and 500 through recent experience), it should be obvious that, except in rare circumstances, these agents will act like the SOP agents. Thus, in general, there should not be any organizational differences based on training.

These analyses were run for organizations composed entirely of either experiential agents or procedural agents. The results indicate that there is no significant difference between organizations of experiential or SOP agents. Thus, the results for these two cases are pooled. The lack of difference in these two cases indicates that all agents are acting as though they are fully trained and the rate of organizational strategic adaptation examined is slow enough that agents have time to learn to "work within the system" before it changes again. Thus, there are no "memory" effects. Performance differentials are primarily a function of strategic adaptation and not individual learning.

5. The infeasibility of optimization

In theory, if the set of strategies considered is sufficient, if there is an optimal design for that environment, if the organization does not change its cost function, and if the environment does not change, then the organization should eventually be able to locate the optimal design. At least, that is the implication of much of organization theory. But how feasible is it that an organization will be able to locate the optimal design? The answer depends on the shape of the performance surface; i.e., the number of performance peaks and valleys and how they are distributed in the space of organizational designs. For example, if the optimal design (a slim peak) is surrounded by a set of extremely sub-optimal designs (a big wide valley), then it is unlikely that the organization will discover it. In contrast, if the optimal design is surrounded by nearly optimal designs, then it is more likely that the organization will locate the design. Organizational designs, however, can be characterized along a large number of dimensions. Even in this simple model, these dimensions include the size of the organization, organizational density, the number of isolates, the number of decision factors overlooked, the number of agents working on each task, the amount of training the agents have, and so forth. The performance surface is thus truly multi-dimensional and cannot be easily graphed. Nevertheless, we can get a feel of what this surface looks like by exploring different pairs of dimensions. Basically, the underlying performance surface is "not well behaved". This is illustrated graphically in figure 2, where the performance surface by size and organizational density is displayed. Additionally, on the bottom plane, the position of all organizations with 85% or better performance is shown. These near-optimal organizations do tend to lie along a ridge

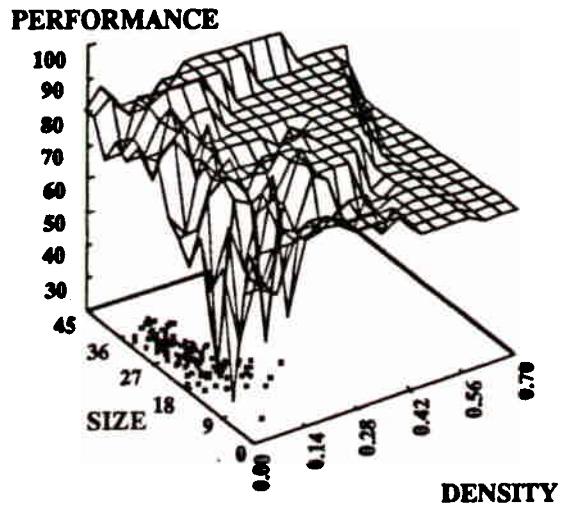

Figure 2. Organizational performance by size and density.

Table 2

High performance plateau: Statistics on organizations with greater than 90% performance.

Variables	Size	Redundancy	Density
Mean	26.46	4.98	0.09
Standard deviation	7.35	1.27	0.04
Minimum	11.00	2.38	0.06
Maximum	38.00	6.78	0.20

(size 18–36, and density 0.07 to 0.14), but known optimal points do not lie on this ridge.

Regardless of the dimensions chosen, the performance features we have just seen appear. That is, there are equiperformance plateaus with sporadic and narrow peaks more or less randomly popping up. The known optimal designs are not surrounded by other designs which exhibit almost as good performance; rather, they are surrounded by valleys of very low performance. For example, one optimal design is at size 9 density 0 and another at size 18 density 0.25. Both of these points are surrounded by low-performing designs or performance valleys. The organizations with better than 90% performance tend to be tightly clustered (see table 2). The results of these analyses suggest that although all organizations might find better designs, they will rarely, if ever, find the optimal design.

Indeed, once the organization is allowed to adapt its design, we find that this is exactly what happens. Initially, all organizations only make the correct decision

50% of the time. Over time, the set of organizations end up normally distributed in terms of performance, with the average performance level around 79% when both modes of adaptation are possible and where none reach 100% performance. Over time, all the organizations tend to gravitate toward the high-performance plateaus; across all organizations, when both modes of adaptation are possible, the final average size is 26.15 and the final average redundancy is 4.01. However, on average, the emergent organizations tend to drop their density too much (final average is 0.06).

6. Successful adaptation

Over time, all organizations change and their performance improves. Some organizations, however, improve more than others. Some organizations are more successful in their adaptation. The most successful organizations tend to be highly flexible; i.e., they hire more, they fire more, they re-assign more agents, and they re-engineer the organization more often. In fact, the most successful organizations make almost twice as many changes (re-assignments plus re-engineerings) than do the least successful (see figure 3). In the most successful organizations, there is a tendency to hire more than to fire (see figure 4), thus retaining more experienced agents. Moreover, the most successful organizations tend, on average, to have a higher ratio of change to their size (number of agents).

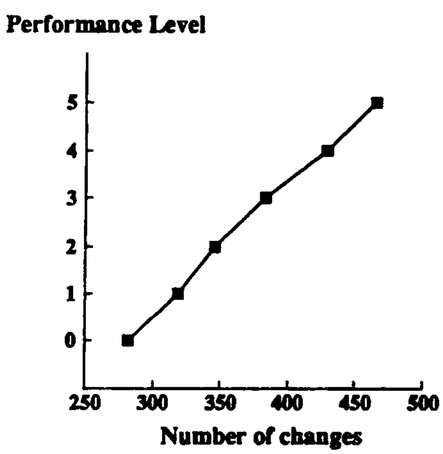

Figure 3. Average number of changes per performance level across all organizations.

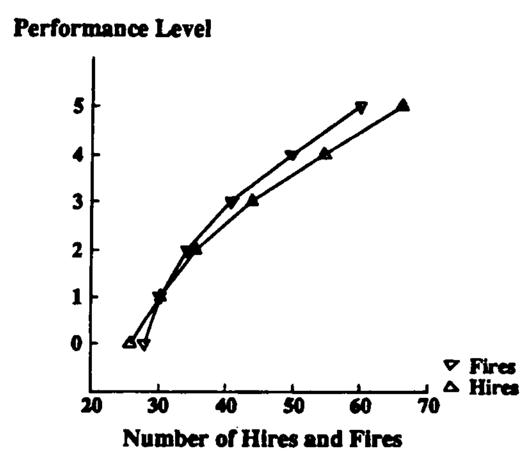

Figure 4. Average number of hires and fires per performance level across all organizations.

Are we seeing a case of success breeding success? All organizations start out with the same initial performance level, the same level of success. However, they do differ initially in terms of structure. And having the "right" initial structure is important for long-term success. This does not mean that organizations that are "poorly structured" can never succeed. Rather, it is possible for them to do so, and the path they take to success may be more torturous. However, it is easier to succeed, and the organization is more likely to succeed, if it is initially well structured. In this way, structure lays the ground work that enables initial successes to breed future successes. Returning to the analogy of a ball on a surface, what we are seeing is that it helps, just a little, to drop the ball in the right place. For example, when both types of adaptation are allowed 15%, the organizations have an initial size, density and redundancy that is more extreme than the mean values in table 2. For these organizations, their final performance is 80.13%. For organizations that were not in this "advantaged" position, the final performance was 78.80%. Thus, there is an advantage to starting out in the right place, but this advantage is quite small.

Dropping the ball in the right place is not sufficient for success. It is also important to roll the ball around a lot. Some organizations with "good" initial designs still end up as low performers. Of the organizations that end up in the lowest performance level, in the long run 16% of them are in this advantaged category in terms of initial size, density and redundancy, whereas, of the organizations that end up in the highest performance category in the long run, only 23% of them are in this advantaged category. Successful organizations not only change faster, they change "smarter". That is, the changes they make are those that are more likely in the long run to be performance enhancing.

Let us look at these results in more detail. Organizations that successfully adapt (end up with the highest performance) have a slight initial advantage in size. For example, in figure 5 we see that organizations with more agents initially do tend to end up more successful. However, although the most successful organizations do not start out the largest; they do end up the largest. In contrast, most organizations start out with comparable density, but the successful organizations are those that most dramatically lower their density (see figure 6). We saw earlier that there was a ridge of near-optimal performers between size 18 to 36 and density 0.7 to 0.14. At each performance level, the average value of the organizations lies within this ridge. Thus, at each level there are some organizations that are in this space. Over time, organizations tend to move (relative to figure 2) up and to the left (increasing size and decreasing density) along this ridge. The organizations, particularly the successful ones, are gravitating toward the spot in the performance landscape where there is the highest density of near-optimal designs. Over time, organizations that end up as low performers end up falling off of this ridge.

Successful organizations are initially more redundant than less successful organizations (see figure 7). This interesting difference, which is only slight initially, is exaggerated over time. The low performers decrease their redundancy slightly. In contrast, over time, high performers alter their design, adding additional agents per task.

Moreover, low performers tend to make strategic errors in the way they change their design. That is, they make changes that actually increase the number of decision

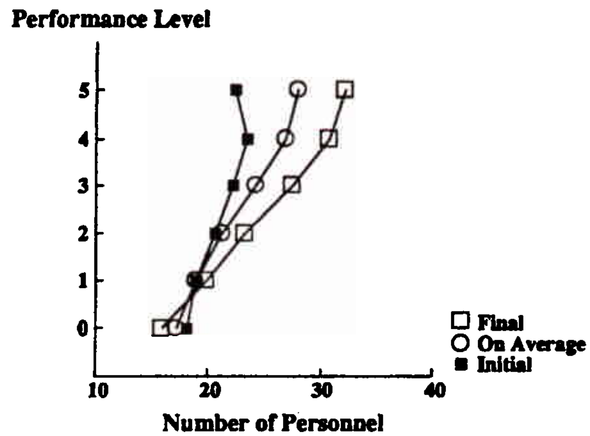

Figure 5. Change in size by performance level.

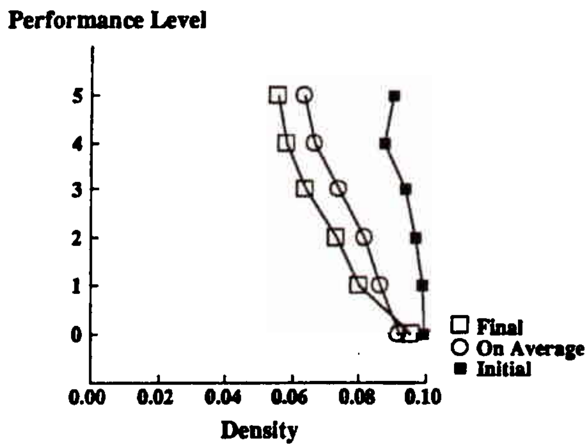

Figure 6. Change in density by performance level.

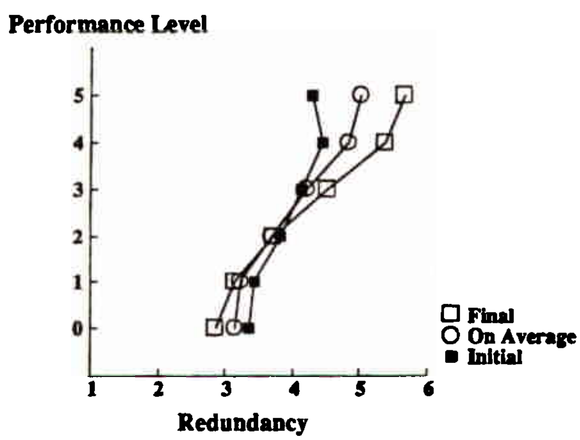

Figure 7. Change in redundancy by performance level.

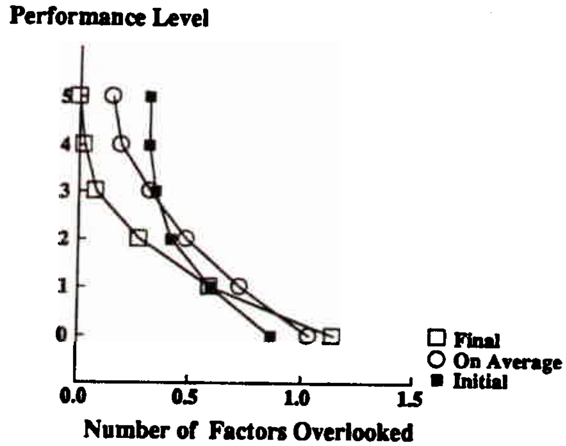

Figure 8. Change in number of decision factors overlooked by performance level.

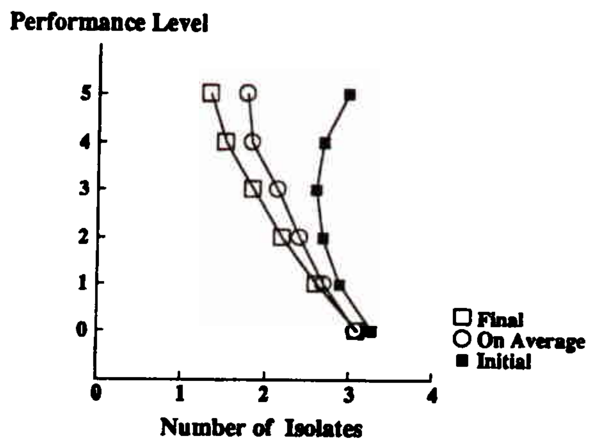

Figure 9. Change in number of isolates by performance level.

factors overlooked (see figure 8), whereas more successful organizations overlook fewer factors. Additionally, there is a more subtle error that unsuccessful organizations are prone to; specifically, they do not get rid of integrate isolates into the organization. Notice that both high-performance and low-performance organizations start out with high numbers of isolates (see figure 9). However, the successful adapters alter their design (sometimes through judicious firing, sometimes through re-assignment, and sometimes through both) so that they reduce the number of isolates.

7. The impact of strategy

Now let us consider the impact of the types of change the organization can employ. For example, what happens if the only strategies employed are to hire or fire agents?

Table 3

Percentage change by level of final performance.

Variable	Final performance level 0	1	2	3	4	5
Size						
Personnel	- 10.90	37.09	45.58	57.04	75.74	72.33
Connections	0.000	0.000	0.000	0.000	0.000	0.000
Both	- 4.15	35.06	58.71	66.49	47.06	60.55
Density						
Personnel	54.56	4.03	- 14.19	- 21.73	- 21.64	- 7.51
Connections	0.000	0.000	0.000	0.000	0.000	0.000
Both	11.55	- 10.30	- 18.90	- 25.99	- 26.64	- 39.86
Redundancy						
Personnel	- 12.26	3.70	13.23	33.89	59.63	61.53
Connections	0.86	- 0.22	0.30	- 0.28	- 0.28	0.000
Both	- 17.26	3.71	17.41	33.88	33.64	35.25
Isolates						
Personnel	2.19	22.69	4.83	- 9.25	- 25.52	- 28.27
Connections	9.64	9.47	8.68	8.96	7.36	0.000
Both	41.42	18.89	6.65	- 4.49	- 24.57	- 19.62
Factors overlooked						
Personnel	59.46	32.21	- 1.10	- 17.04	- 23.10	- 15.56
Connections	5.94	1.20	2.56	- 1.88	- 2.61	0.000
Both	81.39	20.37	- 2.54	- 17.97	- 11.11	- 23.08

Alternatively, what is the impact of only re-assigning agents or re-engineering the organization? What happens when both change strategies are employed? In table 3, the percentage change in various factors under each of these three change strategies at each level of final performance is shown. When the organization only engages in re-assignment and re-engineering. it is not possible for the number of agents or density to change. When the organization can change both agents and connections, organizations which successfully adapt alter the number of agents less and alter the connections more than do organizations which employ only change in agent strategy. Organizations can substitute judicious re-assignment and re-engineering for staff augmentation to achieve high performance (connection changes). If organizations are willing to be flexible in altering their structure, they need not employ large structures. With respect to redundancy, we see that organizations which do not successfully adapt use re-engineering to decrease their redundancy. Further, organizations that can alter the number of agents, if they are to be successful, add agents so as to increase redundancy.

When organizations only change connections, only three organizations ended up in the most successful category. All three of these organizations have the feature that they neither isolate individuals nor overlook decision factors. Importantly, they also (unlike some other organizations) do not alter their structure so as to increase the number of isolates or factors overlooked. Barring these three cases, we see that when the organization only changes connections typically, successful adaptation requires isolating agents who are not performing well and dramatically decreasing the number of factors overlooked. When the organization can alter the number of agents, isolates are eliminated and there is even more emphasis on making sure factors are not overlooked.

7. Toward a science of organizing

Computational models such as this embody organizational theory and so can act as hypothesis generators. The results from the virtual experiment can be thought of as hypotheses. That is, it is difficult to theorize about complex adaptive processes so we have used the computational model to generate hypotheses. The organization scientist can now test these hypotheses with empirical data.

In presenting the results from the computational model, an attempt was made herein to remain agnostic as to whether the artificial agents in the proposed model are computational analogs of humans, non-humans, or some combination. Consequently, the hypotheses might apply in multiple domains, such as human and non-human organizations. How might the results we have just seen be interpreted differently for humans and non-humans? Would the predictions in these two cases be different? Following are illustrations of the types of hypotheses that would be consistent with the results previously presented. These hypotheses go a bit beyond the model, but they indicate the kinds of hypotheses that would make sense given this analysis.

Let us begin by assuming that the agents are humans. The foregoing results suggest that individuals who join large organizations are more likely to remain with the organization longer and to see the organization grow around them. As organizations become more successful, the individuals within the organization will come to interact with other individuals within the organization less (density decrease). This should be due, in part, to an increase in the size of the organization. Nevertheless, this decreased interaction may result in individual's feeling increasingly severed from the organization and feeling that the organization is becoming increasingly bureaucratized. In such cases, one might hear statements like "I used to know everybody and their families, I hardly know most of the new people" or "this organization no longer cares about people". These results also suggest that in organizations that are becoming increasingly successful, people might find themselves increasingly surrounded by others doing the same or similar tasks, so that no individual is as critical as he or she was in the earlier, less successful, days. This might lead to an increasing feeling that they are not particularly important to the organization, or it might lead to the feeling of relief that they now have someone to share the job with. Despite this feeling of disconnectedness, these results are also suggesting that the likelihood of an individual actually being isolated is higher in unsuccessful organizations and increases as organizations become increasingly unsuccessful. Thus, while individuals in increasingly successful organizations may feel that they are becoming less involved in the organization, their actual likelihood of being isolated is lower. In contrast, in organizations that are becoming increasingly unsuccessful, people might find themselves increasingly relied on to do more tasks and with fewer others doing the same thing. In this case, we would see individuals becoming critical to the organization's performance to the point that, ultimately, the loss of one individual may be catastrophic. These results also suggest that organizations that are increasingly successful will come to overlook fewer decision factors. However, this result may be due to cleverness in the adaptation strategy chosen and correctly timed risks, rather than a particular ability on the part of the CEO (or corporate board) to correctly anticipate the future. These results suggest that organizational learning resides not just in the minds of the personnel, but also in the connections among them and in the connection between people and tasks. This is different to saying that organizational learning resides in rules and procedures. Rather, it is saying that the structure or design of the organization itself holds knowledge.

Now, how do we interpret these same findings when the agents are non-humans. Let us imagine, for example, that the agents are webbots [10], artificial agents on the web who have some type of information processing capability, accepting some input and passing some output that can potentially come from or be used by another webbot.

For artificial agents such as webbots, the foregoing results suggest that adding more webbbots to an integrated collection of these artificial agents will increase performance. At one level, this hypothesis is hardly surprising as additional webbots might add additional functionality. Importantly, however, the model would also predict that additional copies of the same webbots would increase performance. After all, the agents in the foregoing model are basically identical, differing only in their experience. These results also suggest that, to improve performance, the interactions among webbots should be structured so as to minimize interaction. That is, a defined set of linkages (that can change) is more advantageous than simply letting all webbots send all output to all other webbots. Further, collections of webbots should become more successful if they restructure these connections to overlook fewer factors. Such restructuring might give the impression of increased intelligence, even without any additional intelligence being given to any of the separate webbots. Importantly, the intelligence here resides not in the webbots but in the connections between them. These results also suggest that to increase performance, collections of webbots need to integrate all webbots – create a reporting structure that actually links them – rather than assume users will know on their own how to transform and move the output from one webbot to become input to another webbot.

The foregoing discussion has briefly suggested how the results of the model might play out in two distinct venues. Underlying this discussion are two important issues. First, can there be a science of organizing? Clearly the predictions for organizations of humans and organizations of non-humans are similar. Imagine for the moment that these predictions held up in both venues. This would suggest that there are general principles of organizing. In this case, an important research endeavor would be to distinguish when an organization of agents is acting as a human organization and not as an organization of non-human agents. The second, and closely related issue, is how accurate does the model of the agent need to be to adequately model human organizations? Earlier research using CORP has suggested that at a macro organizational level, the agents may not need to be very accurate at all. Further research, however, is needed to see whether there are certain tasks, or certain levels of tasks, or certain types of performance, where more cognitively and physically accurate models of humans are needed. Being precise about such items would help set the boundary between a science of organizing and a science of human organization.

Acknowledgement

The author wishes to thank Rangaraj Ramanujam for his helpful comments on an earlier version of this paper.

References

[1] G. Allison, *Essence of Decision*, Little Brown, Boston, MA, 1971.
[2] R. Axelrod and D. Dion, The further evolution of cooperation, Science 242(1988)1385–1390.
[3] R.M. Axelrod, *Structure of Decision: The Cognitive Maps of Political Elites*, Princeton University Press, Princeton, NJ, 1976.
[4] R.M. Axelrod, The evolution of strategies in the iterated prisoner's dilemma, in: *Genetic Algorithms and Simulated Annealing*, L. Davis, ed., Pitman, London, 1987, pp. 32–41.
[5] J. Bruderl and R. Schussler, Organizational mortality: The liabilities of newness and adolescence, Administrative Science Quarterly 35(1990)530–547.
[6] R.R. Bush and F. Mosteller, *Stochastic Models for Learning*, Wiley, New York, 1955.
[7] R. Butler, The evolution of the civil service – A progress report, Public Administration 71(1993) 395–406.
[8] K.M. Carley and A. Newell, The nature of the social agent, Journal of Mathematical Sociology 19(1994)221–262
[9] K.M. Carley and Z. Lin, Organizational designs suited to high performance under stress, IEEE Transactions on Systems, Man, and Cybernetics 25(1995)221–230.
[10] K.M. Carley and M.J. Prietula, WebBots, trust, and organizational science, Working Paper, Social and Decision Sciences, Carnegie Mellon University, 1996.
[11] K.M. Carley and D. Svoboda, Modeling organizational adaptation as a simulated annealing process, Sociological Methods and Research, to appear.
[12] K.M. Carley and Z. Lin, A theoretical study of organizational performance under information distortion, Management Science, to appear.
[13] K.M. Carley, A comparison of artificial and human organizations, Journal of Economic Behavior and Organization, to appear.
[14] K.M. Carley, Organizational learning and personnel turnover, Organization Science 3(1992)20–46.
[15] C. Castelfranchi and E. Werner (eds.), *Artificial Social Systems: 4th European Workshop on Modeling Autonomous Agents in a Multi-Agent World, MAAMAW '92*, S. Martino al Cimino, Italy, 1992.

[16] M.D. Cohen and J.G. March, *Leadership and Ambiguity: The American College President*, McGraw-Hill, New York, 1974.
[17] K. Crowston, Evolving novel organizational forms, in: *Computational Organization Theory*, K.M. Carley and M.J. Prietula, eds., LEA, Hillsdale, NJ, 1994, pp. 19–38.
[18] K. Crowston, An approach to evolving novel organizational forms, Computational and Mathematical Organization Theory, to appear.
[19] P.J. DiMaggio and W.W. Powell, The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields, American Sociological Review 48(1983)147–160.
[20] E.H. Durfee and T.A. Montgomery, Coordination as distributed search in a hierarchical behavior space, IEEE Transactions on Systems, Man, and Cybernetics 21(1991)1363–1378.
[21] R.G. Eccles and D.B. Crane, *Doing Deals: Investment Banks at Work*, Harvard Business School Press, Boston, MA, 1988.
[22] G.S. Elofson and B.R. Konsynski, Performing organizational learning with machine apprentices, Decision Support Systems 10(1993)109–119.
[23] L. Gasser and A. Majchrzak, ACTION integrates manufacturing strategy, design, and planning, in: *Ergonomics of Hybrid Automated Systems IV*, P. Kidd and W. Karwowski, eds., IOS Press, The Netherlands, 1994, pp. 133–136.
[24] M.T. Hannan and J. Freeman, The population ecology of organizations, American Journal of Sociology 82(1977)929–964.
[25] J.H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, MI, 1975.
[26] J.H. Holland and J. Miller, Artificial adaptive agents in economic theory, American Economic Review, Papers and Proceedings 81(1991)365–370.
[27] J.R. Hollenbeck, D.R. Ilgen, D.J. Sego, J. Hedlund, D.A. Major and J. Phillips, The multi-level theory of team decision making: Decision performance in teams incorporating distributed expertise, Journal of Applied Psychology 80(1995)292–316.
[28] J.R. Hollenbeck, D.R. Ilgen, D. Tuttle and D.J. Sego, Team performance on monitoring tasks: An examination of decision errors in contexts requiring sustained attention, Journal of Applied Psychology 80(1995)685–696.
[29] E. Hutchins, The technology of team navigation, in: *Intellectual Teamwork*, J. Galegher, R. Kraut and C. Egido, eds., LEA, Hillsdale, NJ, 1990, pp. 191–220.
[30] E. Hutchins, The social organization of distributed cognition, in: *Perspectives on Socially Shared Cognition*, L.B. Resnick, J.M. Levine and S.D. Teasley, eds., American Psychological Association, Washington, DC, 1991, pp. 238–307.
[31] T. Ishida, L. Gasser and M. Yokoo, Organization self-design of distributed production systems, IEEE Transactions on Knowledge and Data Engineering 4(1992)123–134.
[32] R.H. Kilmann and T.J. Covin (eds), *Corporate Transformation: Revitalizing Organizations for a Competitive World*, Vol 1, The Jossey-Bass Management Series, Jossey-Bass, San Francisco, CA, 1988.
[33] D.H. Kim, The link between individual learning and organizational learning, Sloan Management Review (Fall 1993)37–50.
[34] S. Kirkpatrick, C.D. Gelatt and M.P. Vecchi, Optimization by simulated annealing, Science 220 (1983)671–680.
[35] T.L. Lant, Computer simulations of organizations as experimental learning systems: Implications for organization theory, in: *Computational Organization Theory*, K.M. Carley and M.J. Prietula, eds., LEA, Hillsdale, NJ, 1994, pp. 195–216.
[36] P.R. Lawrence and J.W. Lorsch, Organization and environment: Managing differentiation and integration, Graduate School of Business Administration, Harvard University, Boston, 1967.
[37] R.E. Levitt, G.P. Cohen, J.C. Kunz, C.I. Nass, T. Christiansen and Y. Jin, A theoretical evaluation of measures of organizational design: Interrelationship and performance predictability, in: *Computational Organization Theory*, K.M. Carley and M.J. Prietula, eds., Lawrence Erlbaum, Hillsdale, NJ, 1994, pp. 1–18.

[38] B. Levitt and J. March, Organizational learning, Annual Review of Sociology 14(1988)319–340.
[39] Z. Lin and K.M. Carley, Organizational response: The cost performance tradeoff, Management Science, to appear.
[40] Z. Lin, Organizational performance – Theory and reality, Doctoral Disserrtation, H.J. Heinz III School of Public Policy and Management, Carnegie Mellon University, 1994, unpublished.
[41] J.G. March, Footnotes to organizational change, Administrative Science Quarterly 26(1981) 563–577.
[42] J.G. March and H. Simon, *Organizations*, Wiley, New York, 1958.
[43] K. Mihavics and A. Ouksel, Learning to align organizational design and data, Computational and Mathematical Organization Theory 1(1996)143–155.
[44] M.L. Minsky, *The Society of Mind*, Simon and Schuster, New York, 1988.
[45] A. Ouksel and K. Mihavics, Organizational structure and information processing costs, Technical Report IDS-95-1, College of Business Administration, The University of Illinois at Chicago, 1995.
[46] J.F. Padgett, The emergence of simple ecologies of skill: A hypercycle approach to economic organization, in: *The Economy as a Complex Evolving System*, B. Arthur, S. Durlauf and D. Lane, eds., Santa Fe Institute Studies in the Sciences of Complexity, to appear.
[47] A. Pete, K.R. Pattipati and D.L. Kleinman, Distributed detection in teams with partial information: A normative descriptive model, IEEE Transactions on Systems, Man, and Cybernetics 23(1993) 1626–1648.
[48] M.J. Prietula and K.M. Carley, Computational organization theory: Autonomous agents and emergent behavior, Journal of Organizational Computing 41(1994)41–83.
[49] E. Romanelli, The evolution of new organizational forms, Annual Review of Sociology 17(1991) 79–103.
[50] R.A. Rutenbar, Simulated annealing algorithms: An overview, IEEE Circuits and Devices Magazine 5(1989)12–26.
[51] H.A. Simon, Decision-making and administrative organization, Public Administration Review 4 (1944)16–31.
[52] H. Simon, A behavioral model of rational choice, Quarterly Journal of Economics 69(1955)99–118.
[53] H. Simon, Rational choice and the structure of the environment, Psychological Review 63(1956) 129–138.
[54] B.M. Staw, L.E. Sanderlands and J.E. Dutton, Threat-rigidity effects in organizational behavior: A multilevel analysis, Administrative Science Quarterly 26(1981)501–524.
[55] A. Stinchcombe, Organization-creating organizations, Transactions 2(1965)34–35.
[56] A. Stinchcombe, Social structure and organizations, in: *Handbook of Organizations*, J.G. March, ed., Rand McNally, Chicago, 1965, pp. 153–193.
[57] Z. Tang, K.R. Pattipati and D.L. Kleinman, A distributed *M*-ary hypothesis testing problem with correlated observations, IEEE Transactions on Automatic Control 37(1992)1042–1046.
[58] H. Verhagen and M. Masuch, TASCCS: A synthesis of double-AISS and plural-SOAR, in: *Computational Organization Theory*, K.M. Carley and M.J. Prietula, eds., LEA, Hillsdale, NJ, 1994, pp. 39–54.