Adaptive Organizations and Emergent Forms

Kathleen M. Carley

Carnegie Mellon

Abstract

Over time organizations change and coordinate personnel in new ways. Such changes may be precipitated by actual or anticipated changes in personnel, the environment, technologies, legislation, or the top management team. This adaptation is constrained and not all forms of coordination are feasible. Since organizations are inherently computational entities insight is gained by examining the adaptation of organizations using intelligent artificial agents. Using ORGAHEAD, a multi-agent model of organizational behavior, a series of virtual experiments were run to examine issues of organizational adaptation. Results suggest the concurrent occurrence of experiential learning and structural learning generates within the organization the ability to learn meta-change strategies which can be either adaptive or maladaptive. Such meta-change strategies effectively lock organizations into divergent paths of behavior which produce heterogeneity of form across the population of organizations. Organizational performance and form depend on a complex of array of factors including environmental change, experiential and structural learning, and the emergence of institutionalized strategies.

Adaptive Organizations and Emergent Forms

The thought of an organization conjures up images of people working to accomplish tasks, managers coordinating personnel, mergers, layoffs, hiring practices, and so forth. Organizational performance is often argued to be linked to the process of production and the technologies involved. Yet, many modern organizations trade in knowledge not just goods, in producing services and information and not just physical devices. As such issues of knowledge, learning, adaptivity become paramount. Modern organizations are essentially knowledge consumers and producers. The ability to produce and use knowledge depends, at least in part, on the personnel in the organization. Yet, individuals change jobs and frequently move between organizations.

The relationship between individual learning (micro-behavior) and organizational learning and adaptation (macro-behavior) is the subject of much debate in the organizational sciences [15]. Examining organizations from a computational socio-knowledge perspective; however, clarifies aspects of this relationship and provides us with important insights and predictions vis organizational behaviors [6]. Herein, this perspective is described. A computational model that embodies this perspective, ORGAHEAD, is briefly presented. Some of the predictions and insights afforded by the model are then discussed.

COMPUTATIONAL SOCIO-KNOWLEDGE PERSPECTIVE

Recent advances in social networks, cognitive sciences, computer science, and organization theory have led to a new perspective on organizations that takes into account the computational nature of organizations and the underlying social and knowledge networks [16]. At the heart of this perspective is the argument that organizations are complex, computational and adaptive [7]. They are also synthetic agents composed of other complex, computational and adaptive agents constrained and enabled by their position in a social and knowledge web of affiliations linking agents, knowledge, and tasks. Meaningful insights into organizational behavior can be gained through the use of computational models.

We expect this to be true, whether the organization is a collection of people, a collection of artificial agents, or a combination of the two. The only difference, is that different agents will have different cognitive and communicative abilities, different capabilities for acquiring, processing, storing, retrieving, communicating information. The capabilities of the agents will define what types of “social” behaviors emerge [10]. This computational socio-knowledge view of organizations will be elaborated. To orient the reader, however, the basic precepts are summarized in Table 1.

Table 1. Characteristics of the Computational Socio-Knowledge Perspective
Organization as complex, computational and adaptive
Knowledge is structural
Organizational knowledge is distributed
Organizations contain networks of multiple agents Agents are human or artificial Networks are enabling and constraining
Organizational action and the action of agents within the organization is embedded in the network
Agents exhibit Structural Rationality
Organizational decision making is search Search may either explore new parts of the knowledge and interaction space or exploit old competencies Search may be either local or global
Constraint based adaptation
Organization exhibits liability of newness

Multiple levels of learning, adaptation, evolution
Permeable boundaries around agents, tasks, design

The basis for arguing that organizations are complex, computational and adaptive agents lies in a structural and "bodiless" view of knowledge. A structural view of knowledge is based on the characterization of knowledge as a complex structure of ideational kernels and the connections among them. This basic notion that knowledge can be broken down into component bits of information and relations is at the core of more knowledge representation schemes whether frames, or rules, or mental models. Synthetic adaptation derives from the structural nature of intelligence. For individual agents, knowledge is the set of ideas or concepts or bits of information known and the relations among them. The relations among information can occur within the mind of an agent or between agents, such as "shared ideas" or the "I know that you know" linkages. Knowledge exists within and between individual agents. Consequently, knowledge must exist within and between any group that contains agents; i.e., knowledge must exist within composite agents. The intelligence of humans is often seen as being composed of two factors – the cognitive architecture and the knowledge. Again, this idea is not particularly new. The new insight is simply that agents can contain agents, that relationships connecting bits of information can extend among agents, and that cognitive architectures can cut across multiple agents. In other words, knowledge exists within and among agents and is not directly linked to the physical body of any one agent.

Within an individual, learning leads to the bits of information being learned and forgotten (nodes added and dropped) and connections among information being learned and forgotten (relations added and dropped). By simple extrapolation, any agent that can add or drop nodes or relations in the knowledge space is learning. Agents, can be composed of other agents. Organizations, for example, can be viewed as composite agents composed of human and artificial agents. When the member agents change, for example, by entering, leaving or moving within the organization, the information nodes are changing (and possibly some relations). When the member agents communicate and learn from each other the relations among information changes at the organizational agent level. In organizations, and indeed in any composite agent, knowledge is distributed. Change in this knowledge is synthetic adaptation. Synthetic adaptation occurs in a composite agent as the member agents change or the connections within and among those agents change. Any entity composed of intelligent, adaptive, and computational agents is also an intelligent, adaptive, and computational agent. In this sense, organization are themselves intelligent, adaptive and computational agents in which learning and knowledge are distributed [5,13].

Due to synthetic adaptation, groups and organizations are more than the simple aggregate of the constituent personnel. Organizations are themselves intelligent, adaptive and computational entities [7]. Consequently, organizations can take action distinct from and not predicted on an aggregation of the actions of member agents. This is a synthesis not an anthropomorphisis. For organizations, their intelligence, adaptiveness, and computational capability result from the detailed, ongoing, interactions among and behavior of the member agents. The synthesis process involves many complex non-linear processes and is not a simple aggregation across member agents. For example, organizational intelligence is not simply the sum or average of the members' intelligence. More complex principles of combination are entailed.

Within organizations, agents, resources and tasks are connected by and embedded in networks of interaction and knowledge [2,11]. The network among agents, the social network, entails all relations by which agents interact, communicate and exchange goods, services, and information. Common instantiations of this network are the communication network (who talks to whom), the authority network (who reports to whom), the friendship network (who is friends with whom) and the advice network (who goes to whom for advice). The network linking agents to bits of information is the knowledge network. This network defines who knows what. We can think of agents, whether people, artificial agents, or organizations, as existing in an interaction-knowledge space defined by these networks. As agents learn, as agents interaction, these networks change dynamically causing agents to move about in this space [12].

For agents, what actions they can take depends on what knowledge they have what knowledge is currently salient. What knowledge is salient is in part a function of the agent's cognitive architecture [10]. The cognitive architecture of humans, makes them at least boundedly rational [10,18]. In fact, they are actually even more limited. Such cognitive limitations create the need for social behavior [10]. Social interaction, and more specifically, the social architecture, the structure of the group, organization or society, also places a limit on agent rationality. Thus, what information is available and salient to the individual is also, in part, a function of the agent's position in the social network. Consequently, all actions are embedded in the social and knowledge networks which curtail the creation, use, and acquisition of information. In this sense, it makes more sense to think of agents not as boundedly rational, but as structurally rational.

For individual agents, decision making is often characterized as search [21]. For the organization, this is also the case. For organizations, there is a relationship between their organizational architecture and its performance. The organization's architecture, its extant structure is typically characterized in terms such as size, span of control, density of connections among personnel, workload, and so on. If we take the set of all organizations, for a particular task, there is a performance surface which characterizes performance for all possible architectures. For the sake of illustration, a possible performance surface linking organizational performance to two structural features – size and span of control is shown in figure 1. This three dimensional figure indicates what performance is possible if the organization's structure is of a particular size and span of density. We can think of the height of a peak as the maximum possible performance achievable with this architectural configuration.

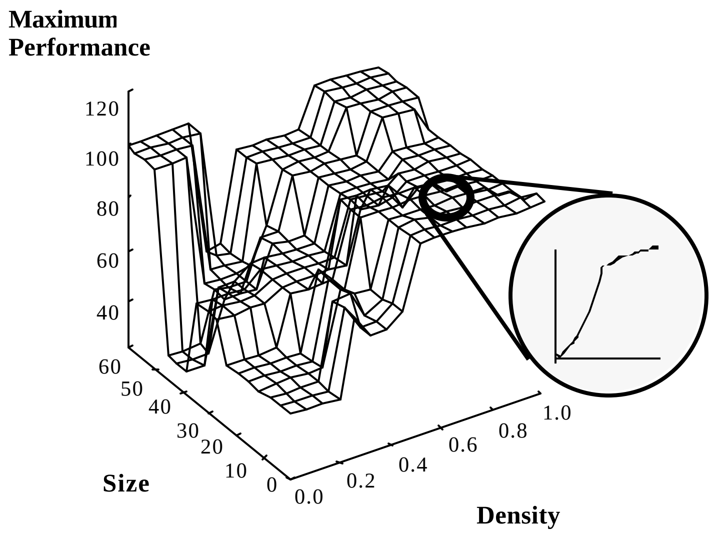

Figure 1. Performance surface for organizations looking at size and density as the relevant structural factors.

The organization, as a synthetic agent, is trying to move through this space searching for the structural form that enables higher performance. The actual search is done by various change agents or processes in the organization, such as the CEO. As structural changes are made, the organization moves about in this space. Structural changes include such activities as hiring new personnel, eliminating divisions, redesigning the organization so that who is reporting to whom changes, and retasking individuals so that who is doing what changes. This search may be done locally – looking only at organizations who are nearby in the performance space and so have similar architecture – or globally – looking at the vast panorama of what architectural forms are possible.

This search may be done in an exploratory manner – looking out at totally new forms – or it may be done in such a way to exploit known competencies – for example, moving along only a single dimension [17]. While organizational theory has examined the potential impact of these search strategies, it is important to note that all the search strategies that humans are capable of are also available to organizations. This does not mean that organizations can search and change in any way they wish.

Rather, the adaptation of agents (human, artificial or organization) is constrained, often by forces beyond the agents immediate control. Laws, financial constraints, technology, etc. limit where in the performance space organizations can move. For example, the tenure system effectively curtails the search through the performance space for universities. In general, the organization will have a set of change strategies and these will determine how it moves through this space.

Such changes in the organization's architecture, because they alter the organization's position in this performance space, may alter the maximum performance that is possible. Changes in the organization's architecture are often characterized as structural learning. Such changes, because they alter the interaction and knowledge networks within the organization, because they may lead to the creation or elimination of knowledge as personnel are moved about, may also effect different levels of actual performance. Since moving the organization from one location to another may eliminate knowledge, or render it less useful, the organization can actually move to a position that structurally enables a higher potential maximum performance but in which, at least initially, the organization performs worse. As an example, installing new technologies and new data bases in organizations, which can be thought of as altering the structure of the organization can temporarily degrade performance as the knowledge people had about how to operate the old technology is no longer valid [1]. When an organization adopts a new architectural form performance may drop and then increase (often following the typical ojival learning curve trajectory) as the individual agents within the organization gain experience in the day to day operation of performing the task in under the new architecture with the maximum attainable value that given by the height of the peak in the performance surface. In figure 1, this is illustrated graphically by the insert in the circle showing the aggregate effects of experiential learning on performance that can occur for organizations that occupy a particular structural position.

Within the organization individual agents learn. One important mode of learning is experiential. For example, as agents engage in quasi repetitive or completely repetitive tasks, they are capable of making generalizations and getting better at the task. However, experiential learning is fragile. Thus, if the organization changes its architecture, the agents within the organization may find that their experience is not valuable. Experiential learning and structural learning clash when the lessons of experiences become invalid as the agent is given a totally new task to do, or a new group of people to interact with.

Initial clashes between structural and experiential learning are more common in new organizations. In new organizations, or in organizations with new top management teams, there is typically a sense of adventure, of risk taking. This leads the organization to explore more organizational architectures, and to engage in relatively frequent structural change. Part of this structural change is inevitable as the organization begins to grow. This frequent structural change inhibits experiential learning. This can in turn decrease performance. A consequence is potentially higher failure rate for new organizations than old organizations. This is known as the liability of newness [23]. Over time, however, most organizations become more staid, less prone to exploration, less willing to take risks.

Experiential and structural learning are two levels of learning within the organization. Since groups, organizations, industries, and so on can be synthetic agents, learning is possible at many levels. Learning involves change in who knows what. Such learning may not be adaptive. That is, agents (human, artificial, organizational) often learn things but that learning does not necessarily serve to improve or maintain performance, that learning does not necessarily advance the agent toward some goal. Adaptation, learning that serves to improve or maintain performance, can also exist at multiple levels. Similarly, evolution, which involves change in the social networks and knowledge networks of sets of agents, can also occur at multiple levels.

The final characterization of organizations, is that even though they can be usefully characterized in terms of collections of agents engaged in tasks, agents and tasks are not inevitably separate entities. The boundaries surrounding agents and tasks are permeable (for a discussion of how this would work in a multi-agent system see [14]). For example, imagine the task of assembling a battery operated toy. For this task, for a human doing the installation there are a large number of subtasks including picking up the toy, a screwdriver, the batteries, unscrewing the battery compartment, placing the batteries inside, screwing back on the battery cover, and setting down the toy and screwdriver. Now imagine a robot that happens to have a built in screwdriver doing this task. In this case the subtasks include picking up the toy, the batteries, unscrewing the battery compartment, placing the batteries inside, screwing back on the battery cover, and setting down the toy. Since the boundaries around agents and tasks are permeable, the design of an organization involves determining what agents, what tasks, and what organizational architecture should be used and not just what architecture should be used. While this view of design may seem obvious to those familiar with developing computational models, it is a rather novel idea to organizational theorists and practitioners who assume that the only agents are humans.

Collectively, this depiction of organizations represents a new perspective on organizations. This perspective, which is drawn from a large number of empirical studies, is consistent with arguments of distributed cognition [13], transactive memory [20,24,25], and the social construction of knowledge [2,5]. Viewing organizations in this way makes it clear that multi-agent models can be meaningfully employed in the development and explication of organization theory.

ORGANIZATIONS AS COORDINATED MULTI-AGENT STRUCTURES

This view of organizations is embodied in the computational framework - ORGAHEAD. ORGAHEAD is a computational framework for examining the behavior of individuals and organizations as they learn, interact, and perform tasks [3,9,12]. Issues relating to organizational learning, adaptation, design, re-structuring, training, agent ability, and so on can be examined using ORGAHEAD. ORGAHEAD has been successfully employed to look at strategic change in small organizations 5-30 agents.

ORGAHEAD characterizes organizations at two levels: operational and strategic. At the operational level, the organization is characterized as a multi-agent system in which the agents can learn and have a position in the organization's architecture which constrains who they communicate/report to, what resources they have access to, and what sub-tasks they are assigned to. At the strategic level, the organization is characterized as a purposive actor; i.e., there is a CEO or executive committee which tries to forecast the future and decides how to change the organization’s structure to meet anticipated changes in the task environment or to improve general performance. As a result of this characterization, ORGAHEAD is comprised of a series of inter-linked models: the operational agent model, the CEO model, and the task model all of which are linked by the organization’s architecture. The architecture is defined in terms of the number of personnel, resources/sub-tasks and the connections among these entities.

ORGAHEAD is multi-agent model of organizational adaptation in which the agents are structurally rational, task oriented and embedded in networks of interaction and knowledge [5]. Organizational performance is affected the CEO's actions at the structural level and the actions of the personnel agents at the operational level. Agents learn through both experience and communication. The interaction and knowledge networks define what role the agents play and impact what the agents learn when. Each agent is modeled as a simple experiential learner capable of doing a classification task, forgetting, gathering information, and communicating opinions. The CEO is a different class of agent and learns from experience and from expectations about the future. The CEO changes the organizational form and so alters these networks. The CEO is modeled as a simulated annealer [12,22] and so has only a limited ability to look ahead, is capable of making erroneous predictions, learns meta-change strategies, alters the organization's architecture, and has as part of its mental model knowledge about who knows what. The CEO can tune the organization in four ways: hiring, firing, redesigning (changing who reports to whom), and retasking (changing who does what). In addition, periodically, the CEO can shake things up by increasing the rate of accepting potentially risky changes [9].

At the operational level the agents are engaged in doing a classification task. There are a large number of possible classification tasks, the one used herein involves classification of a binary string into a binary choice. That is, there is a binary string and the task is to determine whether the pattern of 1's or 0's in that string indicate that the outcome is of type A or B. The simplest version of this task would be to determine whether there are more 1s or 0s in the string. More complex ones, associate certain complex of weighted patterns of bits with an A outcome and others with a B outcome. The string is sufficiently large relative to the information processing capabilities of the agents that no agent can see more than a portion of the string. Thus, the CEO must coordinate the activities of the agents to ensure high accuracy in correctly classifying strings. The complexity of the task environment is a function of string length. A task environment is the set of strings with the same genetic code, i.e., the same algorithmic determinants of outcome and distribution of 1s and 0s in the string.

A large number of studies can be conducted by running virtual experiments within ORGAHEAD. Among the variables that can be manipulated are: initial organizational size and structure, agent cognitive capability, level of pre-training, complexity and variability of task environment, types of allowable structural change, rate of structural change, and rate of change in risk aversion. Using ORGAHEAD a large number of organizational issues can be addressed. Herein, the use of ORGAHEAD to address a number of different organizational issues is illustrated.

EMERGENT FORMS

Although it is known that organizations change, adapt, and alter their architecture there is little understanding of what forms of architecture are likely to emerge over time and under what conditions. We can address this question using ORGAHEAD. Using ORGAHEAD 1000 organizations varying in initial organizational architecture, training, and agent capabilities were simulated for 10,000 time periods (at one task per time period). These organizations were all adapting to a stable environment. That is, the set of tasks faced by the organization were, time period after time period, similar in complexity and potential results. Data on performance, strategic changes, agent training, and the initial and final architecture of the organization were recorded. These data were then statistically analyzed to determine which aspect of the multi-agent model could be described as a linear system.

Results indicate that while it is relatively easy to predict the final size of the organization, it is difficult to predict performance or other structural features such as density with a linear model (see Table 2). Results also indicate that the combination of structural learning (through redesign, retasking, hires and fires) and individual learning (training and memory) as well as initial conditions influence performance and form. Adaptive organizations, i.e., those that exhibit higher sustained performance, learn to be flexible in whom they assign to what (retasking), grow in size, and employ personnel who are more capable (more memory). Organizations learn to be large if they learn not to redesign (change who is reporting to whom) and if their personnel are not over-trained initially. Density, the complexity of the connections among agents, degrades performance and size. Density also does not emerge from a linear process but is correlated with workload.

A naïve interpretation of this linear analysis is that the same behavior is being exhibited by all organizations. However, the low R2, particularly for performance and density, argues the inadequacy of the linear model. In point of fact, what organizational form emerges depends on the history of changes made by the CEO and the initial organizational architecture. Saying this however, gives little insight into the specific forms that do emerge.

Table 2. Standardized Regression Coefficients for Performance, Size and Density

Variables	Performance	Size (E)	Density (E)
Redesign	-0.050	-0.136***	0.027
Retask	0.268***	-0.084*	0.068
Hires	-0.054	1.520***	-0.200
Fires	-0.053	-1.317***	0.151
Training	0.034	-0.115**	0.047
Memory	-0.062*	0.003	0.015
Size (B)	0.052	0.579***	-0.124
Size (E)	0.493***	--------------	-0.306*

Density (B)	-0.017	0.010	0.035***
Density (E)	-0.099***	-0.090***	--------------
Workload (B)	0.083**	0.035*	-0.012
Workload (E)	-0.014	-0.027	0.267***
R2	0.342	0.779	0.249
Significance: * = P < 0.05, ** = P < 0.01, *** = P < 0.001, N = 1000, B=begin, E=end			

To begin to understand what forms are emerging the relative behavior of the top and bottom performing organizations is examined separately. The top performing organizations are defined as the 10% of the 1000 organizations with the highest average performance in the last 500 tasks. Low performers are the 10% with the lowest average performance during this same period. Over time, organizations change in different ways. Moreover, high and low performers, despite comparable initial organizational architectures, end up with dramatically different forms. This is true on every measure of architecture. For example, poor performing organizations tend to shrink in size while high performance organizations tend to increase in size (see Figure 2). Whereas, for density low performance organizations become more dense and high performance organizations become less dense (see Figure 3).

A second virtual experiment was run where 1000 organizations faced with a changing environment

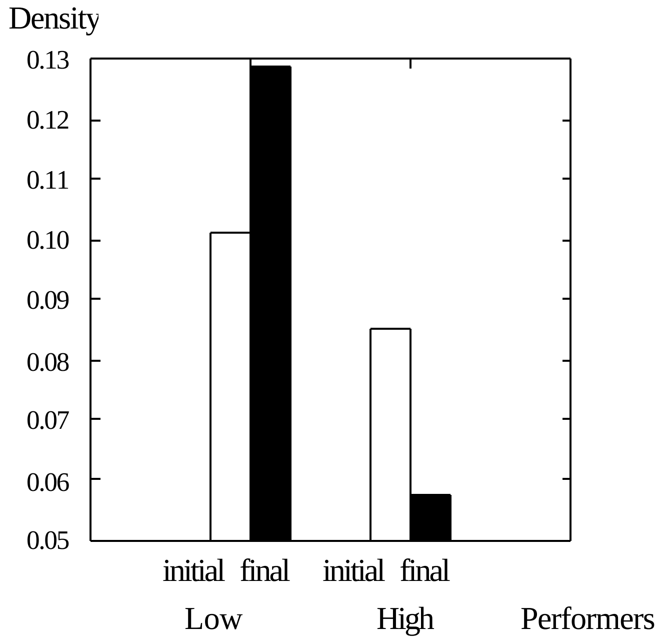

Figure 3. Change in density over time for high and low performance organizations.

were simulated. When the environment is changing, the impact of initial architecture and learning on performance and the form of the organization that emerges is even more complex. Among the findings in this case are the following. Importantly, despite the added complexity of the task facing the organizations, they are still able to learn; however, in this case not all organizations adapt. In fact, the more difficulties that the organization faces simultaneously, the more ways in which the environment changes at once, the greater the decrease in performance. In human organizations, for example, we find that when crises occur and the environment changes the more things that go wrong or change at once the worse the performance. In the case of a volatile environment, maintaining the same level of performance requires a little more adaptation [see also, 4]. In the case of a volatile environment there is a greater need that the meta-learning in the organization be more detailed. Simpler meta-strategies appear less useful. Finally, in such a volatile environment successful organizations tend to exhibit higher density. That is, while many connections among agents are not necessary for adaptation in a stable environment they do appear to be necessary in a volatile environment. This suggests that the social network affords the organization with a type of flexibility that enable s rapid response to, and so higher performance in, volatile task environments. In this case, who you know not what you know affords better performance.

TUNING AND SHAKING

Learning results in organizations that are continually changing. However, only a few of these changes are truly adaptive and so enable the organization to exhibit sustained performance. As noted, in a stable environment less structural change is needed than in a volatile environment to generate the same level of performance. Over time, however, organizations can vary both in the type of structural change in which they engage and the amount of such change. We can think of a sudden increase in the amount of structural change as "shaking" the organization. This is implemented by a sudden increase in the temperature of the system for the annealer (Medeiro cooling schedule [19]). We can think of gradual changes in who is reporting to whom and who is doing what as tuning. With ORGAHEAD we examine the relative impact of tuning and shaking on organizational performance.

In the stable environment, high performance organizations tend to engage in much more tuning than do low performance organizations (see Figure 4). In particular, they tend to bring individual agents in and then expend their effort moving the agents about until the reporting structure stabilizes. Tuning, both redesign and retasking, tend to facilitate adaptation. For high performance organizations, they learn that the optimum type of structural learning is redesign, then retasking, then hiring, then firing. It is important to recognize that had the simulated organization changed its position in the performance space through more aggressive hiring or firing practices it might have ended up in a position with a higher potential maximum performance. However, the organization learns not to make such moves because it learns that structural learning and experiential learning can clash, and that even though such aggressive practices might result in higher potential maximum performance they result in lower actual short term performance. Poor performance firms do not learn this meta-strategy. This simulation result suggests that in human organizations, to the extent that humans have been involved in high performance firms, they will when faced with a choice of changing their organization by tuning, or through downsizing/upsizing will typically choose the tuning change, even if they are told that the architecture that will result from the other changes is optimum from a maximum potential performance standpoint.

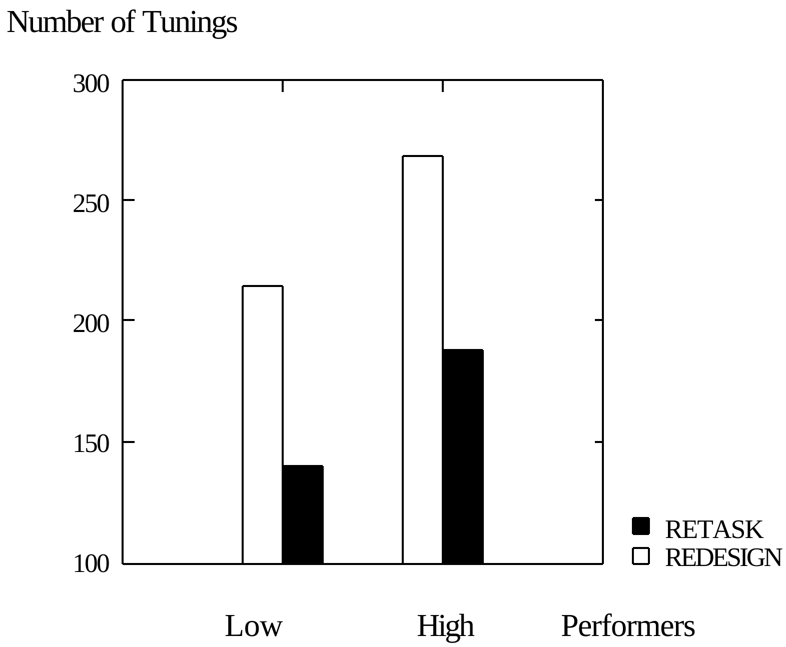

Figure 4. Differences in tuning for high and low performing organizations.

While tuning typically improves organizational performance and leads to greater adaptivity, shaking does not. Shaking involves a temporary increase in risk taking. As such, the organization that shakes itself moves from a position where it only employs change strategy that are expected to improve performance to a position where change strategies that might degrade performance are taken (with increasing likelihood over time due to the annealing process). Earlier work in engineering suggested that such a cooling schedule could actually lead to improved performance for extremely complex response surfaces. In contrast, work with ORGAHEAD suggests that this may not be case. There are two possible reasons for this difference in finding. First, it may be that the mediero schedule works better when the response surface is complex but fixed in time; whereas here, the surface changes slightly in a stochastic fashion with each subtask and dramatically when the task environment is volatile. Second, it may be that in a multi-level model, as exists here with both the multi-agent experiential model and the annealing model combined to characterize the organization, complex cooling schedules are inappropriate. Regardless, from an organization standpoint this work suggests the need for organizations to move with caution if engaging in shaking.

In both the previous sets of virtual experiments, and in these, different behavior appears to be followed by the top and bottom performing organizations. Figures 2 and 3 illustrate that the type of architectural form that emerges is different for low and high performers. Figure 4 illustrates that the way in which low and high performers change is different. At issue, however, is whether this difference is due to the specific set of tasks that are faced by the organizations at the end, or whether low and high performing organizations get trapped into architectural forms and patterns of change that constrain their ability to move in the performance space.

DIVERSIFICATION

Using ORGAHEAD a third virtual experiment was done to look at the tendency to diversification. A set of 100 organizations were simulated for 40,000 time periods in both stable and changing environments. In the changing environment there is a dramatic change in the task at time 20,000. Initial architectural forms, as in the earlier virtual experiments, were chosen randomly. In this case, the over time behavior, and not differences in initial and final states, was observed. Results indicate that the final behavior is due to organizations locking into the wrong strategies and not due to the specific set of tasks in a particular time period.

In a stable environment, over time organizations develop meta-learning procedures which cause them to lock into particular organizational forms. Performance stabilizes although not necessarily at a high level. In a changing environment such meta-learning can be unexpectedly devastating. For example, in Figure 5 the average performance of the 3 most adaptive and the 3 most maladaptive organizations is shown. Oscillations in performance occur as structural changes result in knowledge losses as personnel are laid off, given new tasks or placed in new divisions. Clashes between structural and experiential learning increase diversity in architectural form and performance. Maladaptive organizations lock into strategies of change that are counter-productive (such as oscillating bouts of hiring and firing). In contrast, adaptive organizations lock into strategies that enable continued flexibility (such as repeated tuning through retasking and redesign).

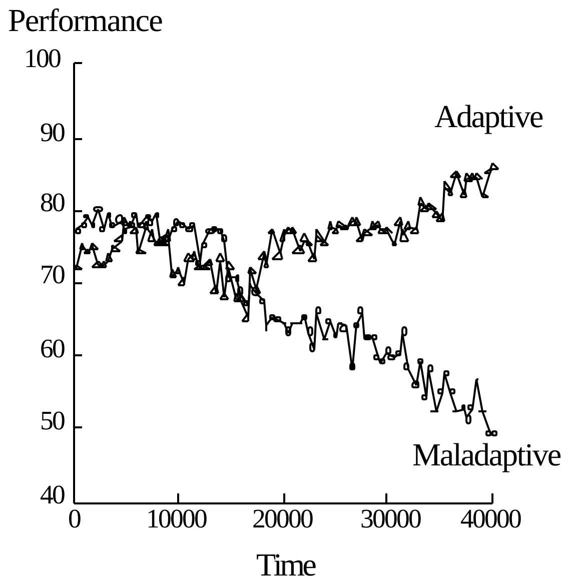

Figure 5. Diversification of organizations over time in terms of performance.

CONCLUSION

The work presented here lies at the intersection of multi-agent analysis and organizational analysis. Advances in both areas are possible by focusing on the development and analysis of veridical models of organizations and the tools to support developing such computational models of organizations. Modeling organizations requires multi-leveled models. In ORGAHEAD, the multi-leveled aspect was provided by combining a set of experiential learning agents with a simulated annealing agent. Other possible multi-leveled models are possible. Additional work is needed to see how the different mechanisms embodied in these multiple models interact, conflict, and support each other. One issue is how sensitive the results are to the specific models. To gain a handle on over time organizational change, constraint based optimization techniques are beneficial at the strategic level. At the operational level, many different types of agent models are possible. For example, in this study, the agents were simple experiential learners. Conceivable, Soar, or other cognitive architectures could have been used to construct the individual agents. Would the results have been different had Soar been used? One previous study suggests that the results might have been slightly different [8]. Whether the differences are systematic and severe enough to lead to different predictions for human organizations is not known.

The difficulty in constructing organizational models, such as ORGAEAD, is that languages that facilitate multi-agent modeling with built in organizational primitives are in, at best, a very rudimentary stage. Future work is needed to develop such languages or frameworks and to ensure that they have appropriate primitives for rapidly constructing, measuring, monitoring, and changing the networks connecting agents and knowledge and not just the primitives for constructing agents. Such languages would advance our understanding of organizations of agents regardless of whether those agents are humans, artificial agents, webbots, robots, or firms.

Thinking about organizations from a computational vantage has led to a new perspective on organizations that enables issues of organizational learning, cognition, and change to be dealt with in a coherent and consistent fashion. The logic described here focuses on the way in which learning mechanisms alter the networks of interaction and knowledge in which agents and tasks are embedded. The implicit argument is that the boundaries around agents, task and resources are to an extent permeable, particularly for composite agents such as teams, groups, and organizations. Let us use the term configuration to denote a specific combination of agents, task, resources and architectural form. The results here go beyond classic findings such as there are multiple architectural forms to achieve any objective, different objectives require different forms, path dependency exists so history and order effects are critical, and overall system behavior is highly non-linear to also suggest that an ecology of learning mechanisms results in organizations engaging in meta-learning about change strategies. Clashes and synergies between experiential and structural learning result in organizations locking into specific change strategies. Such strategies are the meta-learning. Through such meta-learning the norms and procedures developed within the organization become institutionalized. Such meta-learning also leads to the emergence of diversification and heterogeneous behavior at the organizational level. Further, this work suggests that the real issues in organizational design have to due with choosing and moving between configurations, not just choosing and moving between forms.

Herein a new theoretical perspective on organizations was described that draws on issues of computation, social networks and knowledge networks. This perspective argues that organizations are synthetic agents (complex, computational, adaptive, and multi-leveled in their own right) whose behavior is a function of the webs of affiliation linking agents, tasks, resources at both the interaction and knowledge level and the complex, computational and adaptive nature of the member agents. It follows from this perspective that multiple types of learning are possible. Organizational cognition and learning, due to the structural nature of knowledge exists both within the heads of the member agents and in the connections among them. Experiential learning, on the part of member agents, has a direct effect on the knowledge within the member agents heads and structural learning, on the part of the synthetic agent, has a direct effect on the knowledge embedded in the social network. These ideas gain form and specification by their incorporation in a computational model. Herein, such a model, ORGAHEAD, was used to examine organizational change and adaptation. Results indicate that clashes and synergies between structural learning within the synthetic agent and experiential learning within member agents cause the organization to lock into specific change strategies (to engage in meta-learning). Consequently, over time, the population of organizations become more heterogeneous and low and high performing organizations become increasingly different both in form and performance. Truly flexible organizations enable agent learning, adapt their architectural form over time, and adapt how they adapt their form.

Acknowledgement

This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-97-1-0037 and by the National Science Foundation NSF IRI9633 662. The views and conclusions contained in this document are those of the authors 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.

REFERENCES

[1] S. Barley, “The Alignment of Technology and Structure through Roles and Networks,” Administrative Science Quarterly, pp. 61-103, 1990.

[2] K. Carley, “A Theory of Group Stability,” American Sociological Review. Vol. 56.3. pp. 331-354, 1991.

[3] K.M. Carley, “Organizational Adaptation.” Annals of Operations Research. 75: 25-47, 1998.

[4] K.M. Carley, “Organizational Adaptation in Volatile Environments” in C. L. Hulin and D. R. Ilgen (Eds.) Computational Modeling in Organizational Behavior: The Third Scientific Discipline. American Psychological Association: Washington, DC, ch. 3, forthcoming.

[5] K.M. Carley, “On the Evolution of Social and Organizational Networks,” in D. Knoke (Ed) Research in the Sociology of Organizations volume on Networks In and Around Organizations. JAI Press: forthcoming.

[6] K.M. Carley and V. Hill, Structural Change and Learning Within Organizations, CASOS working paper, Carnegie Mellon University, 1999.

[7] K.M. Carley and L. Gasser, “Computational Organization Theory,” in Gerhard Weiss (Ed.) Distributed Artificial Intelligence, MIT Press : Cambridge, MA, Ch. 7, 1999.

[8] K.M. Carley, “A Comparison of Artificial and Human Organizations.” Journal of Economic Behavior and Organization. 31. pp. 175-191, 1996.

[9] K.M. Carley and J. Lee, “Dynamic Organizations: Organizational Adaptation in a Changing Environment,” in Joel Baum (Ed.) Advances in Strategic Management, Disciplinary Roots of Strategic Management Research. JAI Press. Ch. 15. , Vol. 15, pp. 269-297, 1998.

[10] K. Carley and A. Newell, “The Nature of the Social Agent,” Journal of Mathematical Sociology. Vol. 19.4. pp. 221-262, 1994.

[11] K.M. Carley and M. Prietula, “ACTS Theory: Extending the Model of Bounded Rationality,” in K. M. Carley and M. Prietula (Eds.), Computational Organization Theory, Lawrence Earlbaum Associates: Hillsdale, NJ, pp. 55-87, 1994.

[12] K.M. Carley and D.M. Svoboda, “Modeling Organizational Adaptation as a Simulated Annealing Process.” Sociological Methods and Research. Vol. 25.1. pp. 138-168. 1996.

[13] E. Hutchins, Cognition in the Wild, MIT Press: Cambridge, MA, 1995.

[14] D.J. Kaplan, The STAR System: A Unified Multi-Agent Simulation Model of Structure, Task, Agent, and Resource. CASOS working paper, Carnegie Mellon University, 1999.

[15] D.H. Kim, “The Link Between Individual Learning and Organizational Learning,” Sloan Management Review vol. Fall issue. pp. 37-50, 1993.

[16] D. Krackhardt and Kathleen Carley,. “A PECANS Model of Structure in Organizations.”in Proceedings of the International Symposium on Command and Control Research and Technology. Monterey, CA, 1997.

[17] J.G. March, “Exploration and Exploitation in Organizational Learning,” in M.D. Cohen and L.S. Sproull (Eds.) Organizational Learning. Sage: Thousand Oaks, CA, 1996.

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

[19] F. Medeiro, F.V. Fernandez, R. Dominguez-Castro, and A. Rodriguez-Vazquez, “A Statistical Optimization-Based Approach for Automated Sizing of Analog Cells.” In Proceedings of the IEEE Intl. Conf. Computer-Aided Design. Pp. 594-597, November 1994.

[20] R.L. Moreland, L. Argote and R. Krishnan, “Socially shared cognition at work: Transactive memory and group performance.” in J.L. Nye and A.M. Brower. (Eds.) What's social about social cognition? Research on socially shared cognition in small_groups, Sage: Newbury Park, pp. 57-84 1996.

[21] Newell, A., Unified Theories of Cognition, Harvard University Press: Cambridge, MA, 1990.

[22] R.A. Rutenbar, "Simulated Annealing Algorithms: An Overview." IEEE Circuits and Devices Magazine. Vol. 5. pp. 12-26. 1989.

[23] Stinchcombe, A. 1965. "Social Structure and Organizations." Pp. 153-193 in Handbook of Organizations edited by James G. March . Chicago: Rand McNally .

[24] D.M. Wegner, "Transactive memory: A contemporary analysis of the group mind." in B. Mullen and G. R. Goethals (Eds.), Theories of Group Behavior, Springer-Verlag : New York, pp. 185-208, 1987.

[25] D.M. Wegner, "A computer network model of human transactive memory," Social Cognition. Vol. 13.3. pp. 319-339, 1995.

Direct all correspondence to:

Kathleen M. Carley

Dept. of Social and Decision Sciences

Carnegie Mellon University

Pittsburgh, PA 15213

email: kathleen.carley@andrew.cmu.edu

tel: 1-412-268-3225

fax: 1-412-268-6938