Proceedings of the 2004 Winter Simulation Conference
R. G. Ingalls,M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

ON BUILDING AN ORGANIZATIONALLY REALISTIC AGENT-BASED MODEL OF LOCAL INTERACTION AND EMERGENT NETWORK STRUCTURE

James K. Hazy
Brian F. Tivnan

Executive Leadership Program
The George Washington University
Ashburn, VA 20147-2604, U.S.A.

ABSTRACT

We describe research intended to build an agent-based model that is "organizationally realistic." By this we mean that the attributes of the artificial organization of agents conform to empirical results for human organizational systems. We build upon the definitional structure of computational organization theory (Carley and Prietula 1994b) and represent an organization as a network of agents, tasks, resources, and knowledge (Krackhardt and Carley 1998). We do not assume an *a priori* design requirement. Rather, organizational structures are posited to emerge endogenously, the particulars being a key area of study. Agent interactions are governed by local network dynamics, agent-specific rules, and explicit universal constraints (Hazy and Tivnan 2003).

1 COMPUTATIONAL ORGANIZATION THEORY

To provide a basis for the agent model, we build upon the axiomatic definition of an *organization* often used in computational organizational theory and modeling (Carley and Prietula 1994a). This axiomatic base, known as ACTS theory, can be summarized as follows: "organizations are viewed as collections of intelligent agents who are cognitively restricted, task oriented, and socially situated" (p. 56). To look at the artificial organization that results from agent interactions, we adopt a precise description of an organization as a connected network interlinking agents, tasks, resources, and knowledge.

This description, known as the PCANSS or meta-matrix representation (Krackhardt and Carley 1998), has proven to be a useful platform for research, in part because it offers unambiguous structural measurements for testing hypotheses (Carley and Krackhardt 1999; Carley and Ren 2001; Carley, Ren, and Krackhardt 2000). PCANSS takes its name from the sub-matrices that make up the whole or meta-matrix linking each element to all others. Specifically, the sub-matrices are as follows: the Precedence matrix for task-to-task links, the Capabilities matrix for resource-to-task links, the Assignment matrix for agent-to-task links, the Needs matrix for resource-to-task links, the Social matrix for agent-to-agent links, and the Substitutes matrix for resources-to-resources links. In this representation, all knowledge relevant to collective activities is represented as a collection of knowledge nodes that is external to the agents. This controversial aspect of the PCANSS formalism is potentially the most powerful. Relevant knowledge, whether within or outside agent memory, is considered as its own network relating policies, values, instructions, or objectives and may be embodied in manuals, e-mails, databases, or an agent's memory, either as explicit or implicit knowledge (Nonaka 1994).

Agents access knowledge nodes through network connections, either directly or through other agents. They also access resources, are assigned tasks, and communicate with other agents through network connections. First-order links connect a particular node with its neighbors. Links of a neighbor to other nodes are called second-order connections. The network changes as each agent interacts with its local environment over time. For this to occur, agents carry a symbolic representation of their places in the network. As random events, or events initiated by other agents, occur, each agent uses its symbolic representation and its rules to change its local situation. For example, a link to new knowledge may be created (Hazy and Tivnan 2003).

To assess whether emergence occurs in this system, we use Sallach's (2003) definition of *emergence* as being the contributing process of organization to multi-level systems. Sallach asserted that social phenomena emerge from agent (e.g., individual) interactions. This notion provides an ontological basis for our research.

2 ACTION THEORY OF ORGANIZATIONS

To bridge the gap between individual and organization levels of analysis, we adopt an open systems approach (Jackson 2003, von Bertalanffy 1950) and use the general theory of action (Parsons 1951, Schwandt and Marquardt

2000) to evaluate organization-level effects. In particular, we use Schwandt's (1997) Organizational Learning Systems Model (OLSM).

Schwandt (1997) defined organizational learning as "a system of actions, actors, symbols and processes that enables an organization to transform information into valued knowledge which in turn increases its long-run adaptive capacity" (p. 8). Schwandt's OLSM contains four performance and four learning action subsystems, each with its own interchange medium. In this research, we focus on the Environmental Interface subsystem and its medium of interchange, *new information*, and on the Dissemination and Diffusion subsystem and its medium of interchange, *structuration*.

The Environmental Interface subsystem acts as the information-input mechanism for the organizational system (Schwandt 1997, p. 9). This subsystem of actions focuses externally to relate the organization's performance and learning systems to its environment and to develop the means by which it pursues different goals and meets changing environmental conditions. Its interchange medium is new information, which must be found, gathered, and imported into the system.

The Dissemination/Diffusion subsystem moves, transfers, retrieves, and captures information and knowledge for the system. The actions of this subsystem are characterized by their ability to meet the integrating requirements of the other subsystems; they include acts of communication, networking, management, coordination, and implementation, roles supporting the norms associated with the movement of information and knowledge (Schwandt and Marquardt 2000). Structuration (Giddens 1984), the interchange medium, is more than a structure of the social system; it is an integration of organizational structures, roles, norms, objects, and processes that provide this dynamic quality called structuration (Schwandt and Marquardt 2000). Effects of structuration can be considered as changes to the organization network (Hazy, Tivnan, and Schwandt 2004a).

The above context, as the next section describes, provides a lexicon that enables us to define some of the relationships between agent-level phenomena and organization-level effects. Thus, some observed organizational phenomena, certain social structures—for example, the organization's boundary (Hazy, Tivnan, and Schwandt 2003a, 2004a, 2004b)—or social roles (Duong and Grefenstette 2004) can be explicitly shown to be emergent in nature.

3 THE MODEL

To develop an organizationally realistic agent-based Model of Organization, Structural Emergence, and Sustainability (MOSES), we chose to imitate a common coordination situation: interdependent tasks that must all be completed to gain collective reward (Axelrod 1984, 1997; Barnard 1938; Shea and Guzzo 1985). In particular, at the agent level, this situation depicts the classic game theory dilemma of cooperation versus defection in agent choice (Fudenberg and Levine 1998, Nash 1950, von Neumann and Morgenstern 1945); and at the collective level, it depicts the fitness value *to* the many of cooperation *by* the many (Axelrod 1984, 1997; Simon 1990). In an economic or market-based implementation, such as the one described here, this approach amounts to mimicking a production process (Smith 1776/1976) or value chain (Porter 1985) within the collective.

To implement a generic cooperation environment, structurally interdependent and sequential tasks, $i \in N$, were defined (i.e., task i+1 depended for input upon the resources output by task i). Thus, resources moved through a value chain of N tasks from raw materials to final product.

Resources were defined such that each task required a resource as input (i.e., a PCANSS needs connection). Upon the completion of each task, a resource was output that was either used as an input for another task or, if it represented the final product, was immediately converted into energy to be distributed among selected cooperating agents according to the payoff function described in the following section.

Agents were defined as action catalysts within the system, interacting dynamically with their local network connections to change their local connection environment (Hazy and Tivnan 2003). Tasks and resources do not interact spontaneously; rather, an appropriately connected agent must be present to perform the task with its respective input resource. When an agent is connected to a task (i.e., a PCANSS assignment) and to the appropriate resources (i.e., a PCANSS capability), and when the task is connected to the resource as its input (i.e., a PCANSS needs connection), then a reaction occurs: the task is executed, and resources are transformed to move along the value chain. Figure 3 shows these value chain relationships.

Knowledge was defined as a specialized resource, one that determines the efficiency of the catalyzed reaction—an agent connection (also a PCANSS capability) to a higher-value knowledge node leads to more efficient transformation along the value chain. This approach is implemented as greater value in the final product and is thus, ultimately, a higher payoff for distribution among cooperating agents. Higher-value knowledge is implemented in the model as knowledge based upon new information introduced into the environment more recently.

New agent-to-knowledge connection, that is, learning, occurs through agent-to-agent interaction. If the agent's method matrix—its symbolic representation of its environment (Hazy and Tivnan 2003)—is such that it has visibility through an interacting agent to that agent's knowledge connections, then the first agent can become connected to the second agent's knowledge. In other words, it can learn. Although the implementation described here assumes that information and knowledge is always exchanged when agents with appropriate methods matrices interact, conceptually there could be a negotiation, using tags, wherein agents build a level of trust or alignment with other agents. Based on this trust, the agent could decide what kind of information—perhaps even false information—is exchanged in an interaction. Future research could use game theory to inform the exchange logic of interactions. When the knowledge gained can be applied to perform tasks more efficiently, this learning process is called *knowledge diffusion* (Schwandt 1997).

The organization's boundary is defined at the agent level. In particular, it is defined as a limitation to an agent's capacity to act as catalyst within the organization. Agents that are in a position to catalyze task and resource transformation interactions—that is, that can engage in coordinated, collective behavior—are considered to be inside the boundary. Those unable to catalyze transformation interactions are said to be outside the organization's boundary. The state of being inside or outside the organization boundary varies over time. Agents that spend time both inside and outside the organization's boundary are called *boundary spanners*. Agents' boundary spanning costs the organization its potential for transformation events (i.e., efficiency declines). However, when new knowledge originates outside the organization, the potential benefits to efficiency can outweigh this cost (Hazy, Tivnan and Schwandt 2004b). Thus, agent-level boundary phenomena can be seen to represent the tradeoff between performance and learning at the organization level, an example of emergent organizational structure (Hazy, Tivnan and Schwandt 2004a, 2004b).

When an agent is outside the organization's boundary and encounters another agent, an outsider agent, the original agent can learn provided it has visibility into the outsider agent's knowledge connections. This is determined by the agent's method matrix and method function and, in theory, may be subject to interaction dilemmas, like the prisoners dilemma described in game theory (Fudenberg and Levine 1998). Because agents outside the organization's boundary cannot perform or catalyze the organization's tasks, learning connections gained outside the boundary do not initially have context within the organization's interdependence structure, its collective objectives (Nonaka 1994). As such, these new connections are referred to as *new information connections* rather than as knowledge connections when the agent is outside the boundary. Thus, agent interactions outside the boundary lead to information transfer events rather than knowledge diffusion events (Hazy, Tivnan and Schwandt 2004b).

Once the information is carried across the boundary of the organization by the boundary-spanning agent, and once further agent interactions occur—now in the context of the interdependent task and reward environment—knowledge diffusion results. Thus, the mechanism whereby boundary-spanning agents gather new information from outside the organization's boundary and import it into the system for diffusion as knowledge is defined at the organization level (Parsons 1951, Schwandt 1997) using only local agent interactions and locally determined agent capabilities and decisions.

The above describes the logic used to develop an agent-based Model of Organization, Structural Emergence, and Sustainability (MOSES). Sustainability results from the evolution of the system. It is determined to a large extent by the nature of the payoff function described in the next section.

4 THE PAYOFF FUNCTION

The payoff function determines the amount of collective resources that the organization gains following the completion of the final task in the value chain (i.e., the final task includes taking the product to market). Recalling that the organization is sustained only by agents' possessing energy to perform tasks for the benefit of the collective, the payoff function determines the resources available for re-capitalization into raw materials and energy to be distributed to agents.

For this analysis, payoff per unit of final product is assumed to be a function of the following three factors: (a) an arbitrary maximum payoff, (b) the currency of the knowledge used to produce the product (i.e., compared to the most current knowledge), and (c) the turbulence in the environment. We assume that the maximum payoff represents the most value that could be realized when the most up-to-date knowledge is used to create the product. From this maximum payoff, the actual payoff is reduced by a factor that captures both the age of the knowledge used in the production of the particular product and the volatility of the market. If older, outdated knowledge is used, the payoff function adjusts to account for the declining value of aging knowledge. In this analysis, the adjustment factor has three terms: (a) the Knowledge Gap—the difference between the current knowledge generation and the generation of the knowledge actually applied in production, (b) the Turbulence Factor—the number of time steps between knowledge generations (i.e., market volatility), and (c) the Maximum Payoff.

Therefore, for each time step in which the organization produces a final product, it is converted into energy units that are subsequently divided equally between the member agents of the organization identified to receive reward (i.e., the organization maintains an egalitarian reward system) (Axtell 1999). Consistent with the above description, the payoff occurs according to the following equation:

$$Payoff = Maximum\ Payoff * \left\{ 1 - (Knowledge\,Gap)\frac{Turbulence\ Factor}{Maximum\ Payoff} \right\}$$

where the Payoff is assumed to equal a Knowledge Gap adjustment downward from a Maximum Payoff. The effect is that more turbulent environments are more forgiving of knowledge gaps. Of course, knowledge gaps develop more quickly and are more prevalent in turbulent environments. This relationship represents the idea that an extremely high

Maximum Payoff would tend to make an organization less vulnerable to change in the environment. An organization with significant intellectual property protection, for example, would have higher margins and therefore would tend to be less vulnerable to knowledge decay in the short run, while in the long run the same organization needs to maintain an awareness of changing trends in its market (i.e., maintaining its Requisite Variety (Ashby 1956)).

Although this payoff function is used throughout this study, variations in its structure may prove of interest in future research.

5 IMPLEMENTATION DETAIL

The model was coded in JAVA, an object-oriented programming language ideally suited for agent-based modeling. The reader will note that a complete JAVA program of MOSES is available from the authors upon request. It was implemented as a grid-world artificial society (Epstein and Axtell 1996), which was allowed to evolve over discrete time steps. In this instance of MOSES, the agents themselves did not apply a genetic algorithm (Epstein and Axtell 1996, Holland 1975/1992, 1995) that evolved independently of the overall network configuration. However, the agents did include changing knowledge connections, which represented learning in the social context. Agents interacted randomly with objects in the grid world, including other agents. The results of all agent and object interactions were processed during each time step, and these values became the initial conditions for the following time step (Hazy and Tivnan 2003).

The *grid world* is the environment in which action occurs. It is defined as a 150-by-150 grid that wraps around both ends (i.e., a torus). A convex and connected 50-by-150 section of the grid, called the Outback, was defined to be outside of the organization's boundary. The edge of the Outback was the organization's boundary. When agents were in the Outback, they could not perform tasks that transformed resources, although they could learn from other agents.

Located randomly at various positions on the grid, but not in the Outback, were pockets of resources. Both member agents (inside the organization) and non-member agents (in the Outback) moved randomly about this grid at each time step. Agents consumed energy at each time step. Energy was replenished when the final product was completed according to the payoff function and payoff distribution algorithm. If an agent's energy level reached zero, the agent was considered dead and was removed from the game.

Only boundary-spanning member agents could cross the organization's boundary. Thus, most agents were limited to motion either inside or outside the organization's boundary. For simplicity, the term "agent" will be used for member agents. The term "outsider agents" will be used when referring to non-members.

At each time step, all agents moved randomly around the grid. Encounters with resources resulted in their transformation into further refined outputs according to the knowledge level of the agent and the agents' task assignments. Both member and outsider agents interacted amongst themselves and exchanged information and knowledge that *might* have enhanced the ability of the one or the other agent to perform a specific task in future rounds. For simplicity, in the implementation described, the agents were hard coded with a simplistic interaction strategy—to always exchange information—rather than a more complex strategy, such as one used to solve the prisoners dilemma (Fudenberg and Levine 1998, Macy and Skvoretz 1998).

Some number of member agents were designated as boundary spanners: able to cross the organization's boundary and interact with non-organizational actors who possessed new information. All new information was introduced first into these outsider agents. This was how knowledge was refreshed, but the member agent had to return to the organization to pass along the refreshed knowledge in order for it to be useful to the organization. This continued until all agents were dead or otherwise for a user-defined number of steps, because dead agents were not replaced during a run, nor could they change their boundary spanner designation.

6 REPRESENTATIVE RUNS

The following section describes the model behavior using results from three different representative model runs—arbitrarily identified as Model Run 1, Model Run 2, and Model Run 3. Each model run resulted from the same set of parametric conditions: the number of boundary spanners was set to 40 of the 100 organizational members, and the turbulence interval was set to 15 (or three business weeks between the introduction of new information in the environment). Although the parametric conditions remained constant in each of the three representative examples, Figures 1 and 2 illustrate that dramatically different trajectories (or outcomes) resulted in each run.

Figure 1 identifies the number of surviving agents within each model run at any given time, up to a defined suspense date (essentially ten years of life for the virtual organizations). Note that the organizations represented by Model Runs 1 and 2 no longer have any surviving agents by time 233 and 1582, respectively. This result concludes the model run in these two cases and, for our purposes, represents the dissolution of the organization.

Figure 1 also shows a period of stability (i.e., zero agent attrition) occurring early in each of the three model runs. In Model Run 1, the organization appears to have reversed an initial trend of poor performance, but it simply cannot sustain this reversal and rapidly decays. In Model Run 2, the organization initially enjoys an extended period of strong performance (an exploitation focus—March 1991) until it finds itself misaligned with the

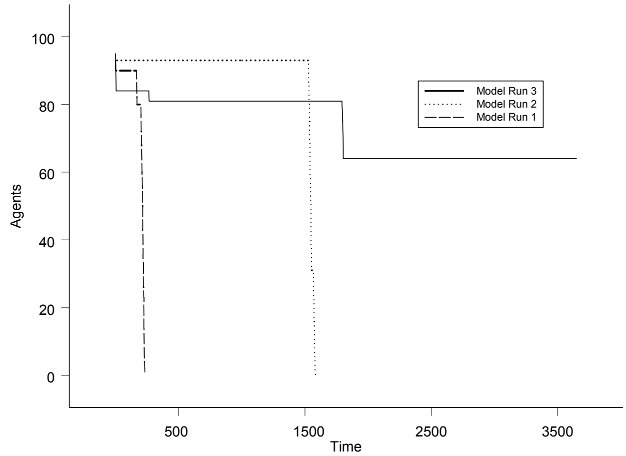

Figure 1: Number of Surviving Agents over Time in Three Representative Model Runs

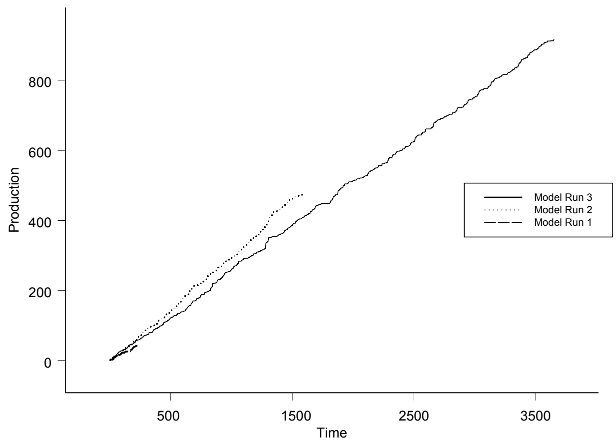

Figure 2: Production over Time in Three Representative Model Runs

changing environmental conditions—a competency trap (Levinthal and March 1993, Levitt and March 1988). In contrast, the organization represented by Model Run 3 appears to overcome similar conditions of initial attrition and an environmental shift.

Figure 2 provides another explanatory perspective for the differentiation between the model runs. Figure 2 illustrates the aggregate organizational production for each of the three organizations. Note that the production rates (i.e., the slope of the respective lines) differ. The organization represented in Model Run 1 begins production almost immediately but at a lower rate, one insufficient to sustain the organization. The organization identified in Model Run 2 maintains production at the highest rate but, due to a lack of boundary spanning, seems to remain unaware that the environmental conditions for its products have changed (e.g., it continues to produce first-generation cellular phones when the market demand has shifted to second-generation digital phones). Alternatively, the organization represented in Model Run 3 experiences a decline in its production rate (Time = 1450) due to changing environmental conditions like those in Model Run 2. This decline in production begins to create a resource-constrained environment within the organization, which eventually leads to agent attrition (Time = 1800). Unlike in Model Run 2, however, the organization in Model Run 3 adapts its product offerings based on the information it gleans from boundary-spanning activity; thus, it resumes its production in such a manner as to sustain the organization in its current state as well as maintain its awareness of environmental conditions.

7 VALIDATION

To gain confidence that the simulation model approximated observed behavior of organizations along the dimensions of interest, we identified stylized facts for comparison. Our hypotheses involved the output level or relative fitness of business organizations producing output by performing a series of interdependent tasks. These tasks were assumed to use resources and belong to a value chain. Tasks used raw materials and converted them into work-in-progress inventory and then to final outputs. We therefore used empirically derived and supported theoretical constructs as our stylized facts.

First, as expected, we observed the production of output along a series of contingent tasks, none of which would produce the final product. Thus, we replicated the value chain—a stylized fact set that we sought to simulate as partial validation of our model (Porter 1980, 1985; Sterman 2000). Within the simulated environment, resources were transformed at various stages of value creation by the action of the agents with the appropriate task assignment and knowledge. Agents consumed energy with each step, and the collective's energy reserve was replenished for agents only when the collective goal was produced—precisely the outcome one would expect in a value chain (Porter 1985). Failure to continually achieve this collective goal—that is, to perform all of the tasks along the value chain—led to the death of individual agents and, eventually, to the end of the collective. As Figure 3 indicates, J independent tasks each transformed one resource, R_j, in the value chain into the next resource, R_{j+1}. When any agent connected to task T_j became connected to resource R_j by random movement, resource R_j was transformed into R_{j+1}.

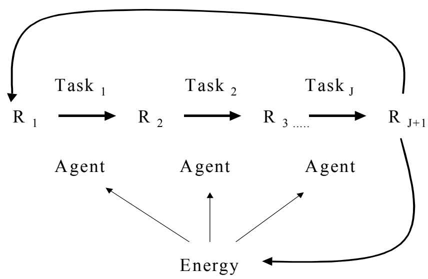

Figure 3: Porter (1985) Value Chain

The above "production process" continued until the completion of final task TJ, wherein a final product, RJ+1, was created and a payoff function was exercised. This payoff function added energy to the appropriate agents and provided new raw resource, R1, to re-initiate the production process. In this way, the collective could sustain itself, and individual agents could survive by benefiting from collective success. Thus, the value chain was effectively simulated.

Second, we sought to simulate collective activity and collective reward, as distinct from individual activity and reward. Thus, we chose to use results from team effectiveness research, such that reward interdependence, task interdependence, and potency could be observed as predictors of successful collective action (Shea and Guzzo 1985). For valuation, results were compared with these studies to demonstrate that the model does indeed simulate collective action: task interdependence, reward interdependence, and collective potency (Lestor, Meglino, and Korsgaard 2002; Shea and Guzzo 1985). Thus, in the base case, no agent could perform all of the tasks itself, and all tasks had to be completed for any reward to be distributed. Also, resources, tasks, and relevant knowledge were available to agents, thereby simulating the collective's ability, or "potency," to execute successfully (Shea and Guzzo 1985). Because no one agent could produce the final good independently, cooperative action and collective success were both necessary for individual survival.

Finally, because the environment was changing, empirical results and organizational theory predicted that without boundary spanners, the organization would eventually fail to respond to the environment and then dissolve. We therefore compare our results at this extreme with the observed decline of organizations that fail to gather new information from the environment (De Vries 1999). Few organizations are able to sustain themselves when low levels of boundary spanning are observed, further supporting the model's validity.

8 COMPUTATIONAL EMPIRICAL RESULTS

To explore the potential usefulness of the above theoretical formulation, we have conducted some initial agent-based modeling computational experiments consistent with the above framework but under simplified assumptions. In this section we briefly describe some results to date.

Agent-based modeling consistent with the previously described network approach was used to study the implications of boundary-spanning activity on organizational learning (Hazy, Tivnan, and Schwandt 2002) and, more generally, the notion of boundary permeability as a construct in agent-based modeling of complex systems (Hazy, Tivnan, and Schwandt 2003b). In these two studies, local agent interactions and the resulting information flows produced the emergent effect of varying social structure, each structure having different fitness in the specific context of the turbulence and complexity of the external environment.

In a third study, the effect of differential rewards to agents on organizational outcomes was studied in the context of agent learning and collective performance (Hazy, Tivnan and Schwandt 2004a). Results of this study showed that when rewards are distributed based upon contribution, either to actual production or to the diffusion of knowledge that informed production, rather than being divided equally among all agents, outcomes improve. Because collective outcomes improve, *an individual agent's survival potential* improves if it participates in production or the diffusion of knowledge—essentially, the result implies that when agents are rewarded for contributions of either exploitation or exploration, collective outcomes improve (March 1991). When rewards are provided to the agents that provided relevant knowledge to other agents, the emergent effect is an increase in organizational performance and sustainability. This offers computational empirical support for (a) individual fitness value of an agent-resident intelligence mechanism that provides visibility into the agent's local network connections and promotes the diffusion of knowledge, and (b) an increased understanding of the emergent relationship between boundary-spanning individuals, organizational learning, and organizational performance.

9 FUTURE RESEARCH DIRECTIONS

The results described above represent first steps toward our goal of an organizationally realistic agent-based model that demonstrates emergent social structure and the dynamics of sustainability in human organizing projects. Future research directions include implementing an emergent task structure and/or payoff function that simulates disruptive technologies and a rugged fitness landscape; allowing new information to develop within the organization as innovation; modeling multiple, interacting organizations, perhaps with overlapping boundaries, and studying the flow of new information between them; adding genetic algorithms and tags to agents, so that they learn and interact according to their own self-interests (e.g., the prisoner's dilemma) (Colomer 1995, Duong and Grefenstette 2004, Fudenberg and Levine 1998, Macy and Skvoretz 1998); and including more complex membership processing (such as trust and reputation), perhaps through tags, that enables exclusionary behaviors, roles, the addition of new agents, and the emergence of an authority hierarchy.

To achieve our long-term goal, many questions remain unanswered: How does collective intentionality emerge, such that emergent agency can be defined at the organization level? How does one represent the synthesis of knowledge into an architecture that includes distinct types of knowledge such as values, norms, and goals, and how does one simulate its effects on local interaction? How are the quality and reliability of information influenced by diffusion within the system (Lawson and Butts 2004)? What are the implications of this effect? How are exploration and exploitation balanced at a system level, and what is the role of leadership (Hazy 2004) in this process? How can computational models help us to explore the complexity of organizational dynamics (Tivnan 2004)? Although these questions are arguably some of the deepest in organization science, we believe that computation methods such as those described here have brought the field to the threshold of understanding. With MOSES, the next steps taken just might lead to the promised land!

ACKNOWLEDGMENTS

The authors have been the recipients of helpful comments and insight from many sources. Of particular note are the suggestions from Dave Schwandt. Additionally, the authors are indebted to their colleagues from both research and practice who participated in the Society for Organizational Learning Research Greenhouse III Conference in Ashburn, Virginia, October 27-29, 2002; the Institute for the Study of Coherence and Emergence, Managing the Complex IV Conference in Fort Myers, Florida, December 7-10, 2002; the North American Association of Computational Social and Organization Science Conference in Pittsburgh, Pennsylvania, June 22-25, 2003; and the International Conference on Complex Systems (ICCS 2004) in Boston, Massachusetts, May 17-21, 2004.

REFERENCES

Ashby, W. R. (1956). *An introduction to cybernetics*. London: Chapman & Hall.

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.

Axelrod, R. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton, NJ: Princeton University Press.

Axtell, R. (1999). The emergence of firms in a population of agents: Local increasing returns, unstable nash equilibria, and power law size distributions. In *Center on social and economic dynamics: Vol. 3. Working papers*. Washington, DC: Brookings Institute.

Barnard, C. (1938). *The functions of the executive*. Cambridge: Harvard Press.

Carley, K. M., & Krackhardt, D. (1999, June). *A typology for C2 measures.* Paper presented at the 1999 International Symposium on Command and Control Research and Technology, Newport, RI.

Carley, K. M., & Prietula, M. J. (1994a). ACTS theory: Extending the model of bounded rationality. *Computational organizational theory*, K. M. Carley & M. J. Prietula, (Eds.), 55-88. Hillsdale, NJ: Lawrence Erlbaum Associates.

Carley, K. M., & Prietula, M. J. (Eds.). (1994b). *Computational organization theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Carley, K. M., & Ren, Y. (2001, June). *Tradeoffs between performance and adaptability for C3I architectures.* Paper presented at the Command and Control Research and Technology Symposium, Annapolis, MD.

Carley, K. M., Ren, Y., & Krackhardt, D. (2000, June 26-28). *Measuring and modeling change in C3I architecture.* Paper presented at the 2000 Command and Control Research and Technology Symposium, Annapolis, MD.

Colomer, J. M. (1995). Leadership games in collective action. *Rationality and Society, 7*(2), 225-246.

De Vries, M. F. R. (1999). High-performance teams: Lessons from the pygmies. *Organizational Dynamics, 27*(3), 66-77.

Duong, D. V., & Grefenstette, J. (2004). An agent-based simulation of social differentiating of roles. *Proceeding from NAACSOS Conference,* K. Carley, ed. Pittsburgh, PA: Carnegie Mellon University.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science form the bottom up*. Cambridge, MA: MIT Press.

Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games*. Cambridge MA: MIT Press.

Giddens, A. (1984). *The constitution of society*. Berkeley CA: University of California Press.

Hazy, J. K. (2004, June 22-25). *Leadership in complex systems: A meta-level information processing capabilities that bias exploration and exploitation.* Paper presented at the NAACSOS, Pittsburgh, PA.

Hazy, J. K., & Tivnan, B. F. (2003). Simulating agent intelligence as local network dynamics and emergent organizational outcomes. *Proceedings of 2003 Winter Simulation Conference,* S. Chick, P. J. Sanchez, D. Ferrin, & D. J. Morrice (Eds.), 1774-1778. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Hazy, J. K., Tivnan, B. F., & Schwandt, D. R. (2002). *Boundary spanning in organizational learning: Preliminary computational explorations.* Paper presented at the Institute for the Study of Convergence and Emergence Managing the Complex IV Conference, Fort Myers, FL.

Hazy, J. K., Tivnan, B. F., & Schwandt, D. R. (2003a). The impact of boundary spanning on organizational learning: Computational explorations. *Emergence, 5*(4), 86-124.

Hazy, J. K., Tivnan, B. F., & Schwandt, D. R. (2003b, May 30–June 2). *Permeable boundaries in organizational learning: Preliminary computational modeling explorations.* Paper presented at the Organizational Learning and Knowledge, Fifth International Conference, Lancaster, UK.

Hazy, J. K., Tivnan, B. F., & Schwandt, D. R. (2004a). Modeling social structure as network effects: Computational evidence that rewarding learning improves performance. *Proceeding from the International Conference on Complex Systems (ICCS2004),* New England Complex Systems Institute (Ed.), 38. Boston, MA: New England Complex Systems Institute.

Hazy, J. K., Tivnan, B. F., & Schwandt, D. R. (2004b, May 20). *Permeable boundaries in organizational learning.* Paper presented at the International Conference on Complex Systems (ICCS2004), Boston, MA.

Holland, J. H. (1975/1992). *Adaptation in natural and artificial systems*. Cambridge, MA: The MIT Press.

Holland, J. H. (1995). *Hidden order: How adaptation builds complexity*. Reading, MA: Perseus Books.

Jackson, M. C. (2003). *Systems thinking: Creative holism for managers*. West Sussex: John Wiley & Sons, Ltd.

Krackhardt, D., & Carley, K. M. (1998, June). *A PCANS model structure in organization.* Paper presented at the 1998 International Symposium on Command and Control Research and Technology, Monterrey, CA.

Lawson, G. C., & Butts, C. T. (2004). Information transmission through human informants: Simulation., *Proceeding of NAACSOS 2004,* K. Carley (Ed.). Pittsburgh, PA: Carnegie Mellon University.

Lestor, S. W., Meglino, B. M., & Korsgaard, M. A. (2002). The antecedents and consequences of group potency: A longitundinal investigation of newly formed work groups. *Academy of Management Journal, 45*(2), 352-368.

Levinthal, D. A., & March, J. G. (1993). The myopia of learning. *Strategic Management Journal, 14*, 95-112.

Levitt, B., & March, J. G. (1988). Organizational learning. *Annual Review of Sociology, 14*, 319-340.

Macy, M. W., & Skvoretz, J. (1998). The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review, 63*, 638-660.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science, 2*(1), 71-87.

Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Science, 36*, 48-49.

Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science, 5*(1), 14-37.

Parsons, T. (1951). *The social system*. New York: The Free Press.

Porter, M. (1980). *Competitive strategy: Techniques for analyzing industries and competitors*. New York: The Free Press.

Porter, M. (1985). *Competitive advantage: Creating and sustaining superior performance*. New York: The Free Press.

Sallach, D. L. (2003). The simulation of emergence: Mysticism, reduction or social complexity? *Proceedings of the 2003 North American Association of Computational Social and Organization Science,* K. Carley (Ed.). Pittsburgh, PA: Carnegie Mellon University.

Schwandt, D. R. (1997). Integrating strategy and organizational learning: A theory of action perspective. *Advances in strategic management, 14,* J. P. Walsh & A. S. Huff (Eds.). Greenwich, CT: JAI Press, Inc.

Schwandt, D. R., & Marquardt, M. J. (2000). *Organizational learning: From world-class theories to global best practices*. Boca Raton: St. Lucie Press.

Shea, G. P., & Guzzo, R. A. (1985). *Leads on the determinants of group effectiveness* (No. WP-533). Philadelphia: The Wharton School, University of Pennsylvania.

Simon, H. A. (1990). A mechanism for social selection and successful altruism. *Science, 250*, 1665-1668.

Smith, A. (1776/1976). *The wealth of nations*. Chicago: University of Chicago Press.

Sterman, J. D. (2000). *Business dynamics: Systems thinking and modeling for a complex world*. Boston: Irwin:McGraw-Hill.

Tivnan, B. F. (2004). Data Farming Coevolutionary Dynamics in RePast. In *Proceedings of the 2004 Winter Simulation Conference,* R. G. Ingalls,M. D. Rossetti, J. S. Smith, and B. A. Peters (Eds.). Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

von Bertalanffy, L. (1950). The theory of open systems in physics and biology. *Science, 3*, 23-29.

von Neumann, J., & Morgenstern, O. (1945). *Theory of games and economic behavior*. Princeton: Princeton University Press.

AUTHOR BIOGRAPHIES

JAMES K. HAZY is a doctoral candidate in the Executive Leadership Program at the George Washington University and a principal in a private capital advisory firm in New Jersey, and he was formerly a financial VP at AT&T and CFO of an Ernst and Young, LLP, business. He earned an M.B.A. with distinction from the Wharton School of University of Pennsylvania and a B.S. degree in mathematics from Haverford College. Jim's research interest is the nature of emergent leadership in complex social systems. His e-mail address is <Jim.Hazy@earthlink.net>.

BRIAN F. TIVNAN is a doctoral candidate in the Executive Leadership Program at the George Washington University and a consultant with the MITRE Corporation. He has a B.S. in mechanical engineering from the University of Vermont and an M.S. in operations research from the Naval Postgraduate School. Prior to attending the George Washington University, Brian served for ten years on active duty in the United States Marine Corps. Brian's research interest addresses the application of complexity theory to organization science. His e-mail address is <BTivnan@mitre.org>.