Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

SIMULATING AGENT INTELLIGENCE AS LOCAL NETWORK DYNAMICS AND EMERGENT ORGANIZATIONAL OUTCOMES

James K. Hazy
Brian F. Tivnan

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

ABSTRACT

We build upon our previous work (Hazy and Tivnan 2003) to represent organizations as a network of agents, tasks, resources and knowledge (Krackhardt and Carley 1998) to explore the emergent effects of agent interactions on organizational outcomes. To do this, we define agents in the context of their position in the network, describe the agent's symbolic representation of its position in the network, and develop a probabilistic function associated with each agent that acts locally to change the network. We conclude with a brief overview of our research in this area to date and the usefulness of this network representation.

1 THEORETICAL BASIS

As a starting point, we accept the axiomatic definition of an organization often used in computational organizational theory and modeling (Carley and Prietula 1994). This axiomatic base can be summarized as: "organizations are viewed as collections of intelligent agents who are cognitively restricted, task oriented, and socially situated" (p.56) and is known as ACTS theory. In addition, we adopt a precise description of an organization as a connected network linking agents, tasks, resources, and knowledge, each type of node is called a different color (Carley, Ren, and Krackhardt 2000).

This description is known as the PCANS or meta-matrix representation (Krackhardt and Carley 1998 and Tsvetovat and Carley 2003). In this representation, all knowledge relevant to collective activities is represented as knowledge nodes external to agents. Agents access knowledge nodes through network connections. They also access resources, are assigned tasks, and communicate with other agents through network connections. First order links connect a node with its neighbors. Links of a neighbor to other nodes are called second order connections.

Although the meta-matrix representation (Krackhardt and Carley 1998) describes the network at a point in time, it does not provide a mechanism to change the network over time. It implicitly assumes homogeneous and inactive agents embedded in a network of connections they cannot change. Other work has developed the ACTS theory by studying the action of agents in a particular task environment (Carley and Prietula 1994). In these studies, the agent's PCANS network is assumed to change, but the mechanism is not explicitly represented (Carley et al. 2000). This work is intended to bridge the gap between these streams of research and describe explicitly, in a manner consistent with both PCANS and ACTS theory, how the action of agents changes the network over time.

2 BRIDGE BETWEEN PCANS AND ACTS THEORY

To bridge the gap between PCANS and ACTS theory, we define an intelligence mechanism within each agent. This mechanism takes as an input its current network state (i.e., its first and second order connections to other agents, to tasks, to resources and to knowledge). Based upon a set of heuristic rules, the agent is then able to change the network connections in its local environment. When one considers the space of all agents, tasks, resources and knowledge, together with all possible connections, then the subset of active connections represents the organization at a point in time.

For this analysis, each agent is assumed to be embedded in a network that defines the collective social situation. As such, each agent can potentially be directly connected to all other agents, all tasks, all resources, and all knowledge components. The links to these nodes are called first order connections. The agent's second order connections are all of the connections of a neighboring node to other nodes that are accessible by the agent through a first order connection. As such, the number of nodes accessible through an agent's first and second order links can become quite large. In reality, however, each agent is connected only to a subset of other agents, tasks, resources and knowledge, and these nodes are likewise connected to a subset of all possibilities. Therefore, an agent's network connections are smaller in practice. For these purposes we define an agent as a node at the nexus of its network connections with other agents, tasks, resources and knowledge. Agent nodes are distinguished from other nodes by having agency, defined to be the active ability to change its network connections. Agent intelligence is the mechanism that enables choices with respect to agency.

3 AN AGENT'S INTERNAL REPRESENTATION OF ITS LOCAL NETWORK

To develop the mathematical model for agent intelligence, we follow and generalize the formulation of Chang and Harrington (2002). We consider a collective consisting of M agents operating as a social system. As a point of departure, rather than simply assuming that individuals engage in an operation decomposable into N tasks as Chang and Harrington do, we assume each agent, $i \varepsilon \{1, 2, 3, \ldots, M\}$ participates in collective activity that can be broken down into S tasks, utilizing R resources, taking advantage of K knowledge components and interacting with $M\text{-}1$ other agents. Like Chang and Harrington (2002), we assume that there are several different methods that can be used by the agent for each task. However, we also assume that there are several methods that can be used with each resource and each component of knowledge. In addition, we assume there are several methods that can be used for each interaction with each other agent. The method chosen by an agent for performing each task, using each resource, interpreting each component of knowledge or engaging each agent interaction is represented by a series of d bits (0 or 1). We define d to equal one plus the number of colors for possible second order connections. Further, because the knowledge resources and tasks are represented as external to the agent, each method is considered in the context of the network connections available to the agent. Specifically, for a particular interaction we assume that each method is fully characterized by the first and second order connections available to the agent in the interaction.

For simplicity, rather than differentiating among interactions with tasks, resources, knowledge and other agents, we will call each of these an interaction. As such there are at most 2^d possible methods for each interaction. In practice the number of methods may be significantly smaller as many combinations are not meaningful. Further, we define total number of possible interactions as

$$N = (M\text{-}1) + S + R + K. \tag{1}$$

Thus, in any time period, t, an agent, i, and its interactions, is fully characterized by a binary N x d matrix, called the agent's "methods matrix". We denote the methods matrix at time t by

$$\underline{z}_i(t) \varepsilon \{0,1\}^{N \times d}. \tag{2}$$

Equation (2) can be decomposed so that

$$\underline{z}_i(t) \varepsilon (\underline{z}_i^1(t),\ldots,\underline{z}_i^N(t)) \varepsilon \{0,1\}^N, \tag{3}$$

and

$$\underline{z}_i^h(t) \varepsilon (z_i^{h,1}(t),\ldots,z_i^{h,d}(t)) \varepsilon \{0,1\}^d. \tag{4}$$

Equations (3) and (4) are agent i's current method in interaction $h \varepsilon \{1,\ldots N\}$. Again, as adapted from Chang and Harrington (2002, p. 7), looking at an example where there are six other agents, six tasks, six resources and six knowledge components (*i.e.*, from Equation (1), there are $N = 24$ possible interactions, and $d = 5$ bits describing the method of interaction). An example of Equation (2), the methods matrix, appears in Figure 1.

Interaction (h):	#1	#2	#3		#24
Interaction method ($\underline{z}_i^h(t)$):	1	0	1	…	1
	1	0	0		1
	0	1	0		1
	1	0	1		0
	1	1	0		1
	←		$\underline{z}_i(t)$		→

Figure 1: An Example of the Methods Matrix

What is shown is that for each one of twenty-four possible interactions, h, whether with a task, a resource, a component of knowledge or another agent, one method of $32 = 2^d = 2^5$ available options is chosen. Given the interactions are completely described by a methods vector of 120 (= 24 x 5) bits, there are 2^{120} possible bit configurations overall (Chang and Harrington 2002: 7-8). Even though in practice this number is much smaller as many combinations are redundant or meaningless, the need for a cognitively restricted agent (Carley and Prietula 1994) to use simplifying heuristics to manage these interactions is apparent.

4 INTERACTION METHODS DEFINED

We define the agent's method for each of its N interactions in the context of the PCANS meta-matrix. As shown in Table 1, for each interaction, the five bits describing the agent's method are defined as follows: the first bit indicates whether a first order connection exists between the focal agent and the relevant object. The second through fifth bits indicate whether the second order connections to other agents, tasks, resources or knowledge nodes can be accessed through the first order connection (i.e., whether the first order connection is strong enough to provide visibility into the next layer of the network).

Table 1: Method Bit Definitions

1	2	3	4	5
First Order Link	Access to Second Order Link to *Agents*	Access to Second Order Link to *Tasks*	Access to Second Order Link to *Resources*	Access to Second Order Link to *Knowledge*
1=Yes 0=No	1=Yes 0=No	1=Yes 0=No	1=Yes 0=No	1=Yes 0=No

As an example, for an interaction with another agent, if bit #1 = 1, there is a connection between the focal agent and another remote agent. If bit #2 = 1, then the focal agent has visibility into the remote agent's connections to other agents. Any heuristic rules governing the agent's network interaction can use information about second order connections. A possible rule would determine when the agent can create a first order connection to augment a second order connection, that is, create a connection with a friend of a friend. However, if bits # 3, 4 and 5 = 0, the focal agent has no visibility into the neighbor's second order links to tasks, resources and knowledge. As such, the agent does not have access to its friend's tasks, resources or knowledge. The method for the first order connection is limited to the social connections of the neighbor, a purely social tie.

This representation embodies information within each agent about its local network. As such it provides the potential for agent action based upon both first and second order connections.

5 MECHANISM OF AGENT INTELLIGENCE: METHOD ENACTMENT FUNCTION

For each agent, for each time step, we define the method enactment function (MEF) as follows: under a set of rules resident in the agent, either stochastic or deterministic, an agent changes its methods of interaction with respect to other agents and/or to tasks, resources or knowledge. These method changes result in changes to the organization's network either as new network connections or as new or transformed resources or knowledge. When the MEF runs during a particular time step, it takes into account current methods, global rules and agent-specific internal rules to determine changes to its methods that will be in effect at the beginning of the next time step. As such, the MEF represents the impact that interactions between an agent and other objects (other agents, resources, tasks, or knowledge) have on that agent's future network connections. The probabilistic outcome of MEF for step t, that is, the subset of the network as updated by the interaction, is defined as that agent's method matrix for time step t+1. During each time step, an agent's method is manifested in local change to the organization's network. The modified network thus becomes the network in effect at the beginning of the next time step.

The limitations to methods defined in this function can be thought of as a simplified (for these purposes) representation of the cognitive limitations inherent in an agent's "mental model" of its local environment (Gavetti and Levinthal 2000). In fact, an agent can enact changes within the environment in which it is connected, but in practice, the agent may be limited to changes allowed within its internal rules, rules that are themselves affected by the agents existing network connections.

6 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 under simplified assumptions. In this section we briefly describe some of the emergent results to date.

We accept Sallach's (2003) definition of emergence, as the contributing process of organization to multi-level systems. Sallach's (2003) assertion that social phenomena emerge from agent (*e.g.*, individual) interactions provides an ontological basis for our research.

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 2003). In these studies involving over 11,000 artificial organizations, certain methods were fixed for all agents to isolate the factors under study and to simplify the analysis. Initially, agents were connected randomly to tasks and tasks to resources. As time elapsed, new "generations of knowledge" were created, connected to various tasks and connected to what were called "outsider agents." Outsider agents remained outside the organization's boundary, that is, they did not participate in collective activities. As agents interacted with one another, including with outsider agents, first order connections were established whenever a second order connection to a new knowledge node became visible to an agent. In this way, new knowledge crossed the boundary (became connected to agents engaged in collective activities) and diffused broadly within the organization. Additionally, these local agent interactions and resulting information flows produce the emergent effect of variation in the social structure of organizations, each structure with different fitness, in the specific context of the turbulence and complexity of it's the external environment (Siggelkow and Rivkin 2003). Stated differently, this process provides an emergent mechanism for variation of organizational forms that depends only upon local interactions and does not depend upon exogenous design.

Results of the above indicated that the level of boundary spanning activity of agents has a non-linear relationship with collective outcomes such as production and number of surviving agents. Figure 2 provides a sample depiction of this non-linear relationship between boundary permeability, defined to be the ratio of information transfer events by agents outside the organization to interaction events by agents inside, and agent survivability (see Hazy, Tivnan and Schwandt 2003 for a detailed discussion of this relationship.)

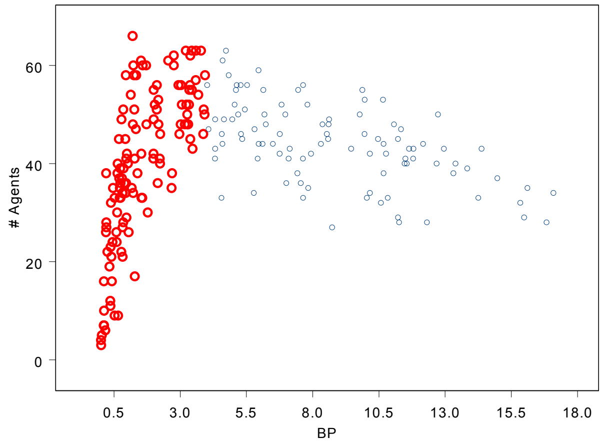

Figure 2: Non-Linear Relationship between Boundary Permeability and Agent Survivability (Hazy, Tivnan and Schwandt 2003)

More boundary permeability increases agent survivability at a diminishing marginal rate until it has little incremental and eventually a negative effect. The specific characteristics of this relationship are dependent upon environmental turbulence, the initial positive effect of boundary spanning being more pronounced with greater turbulence. These computational experiments also demonstrated that when an agent could change its local network, that is, become connected to tasks (i.e., perform) previously not assigned, the emergent outcome that was observed was an increase in organizational production and agent survivability at all levels of boundary spanning.

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 Under review). 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.

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; and the North American Association of Computational Social and Organization Science Conference in Pittsburgh, Pennsylvania, June 22-25, 2003.

REFERENCES

Carley, K. M. and M. J. Prietula. 1994. ACTS Theory: Extending the Model of Bounded Rationality. In *Computational Organizational Theory*, ed. K. M. Carley and M. J. Prietula, 55-88. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Carley, K. M., Y. Ren, and D. Krackhardt. 2000. Measuring and Modeling Change in C3I Architecture. Paper presented at the 2000 Command and Control Research and Technology Symposium, Annapolis, Maryland.

Chang, M.-H. and J. E. Harrington. 2002. Discovery and Diffusion of Knowledge in an Endogenous Network. In *Proceedings of the 2002 CASOS Conference*, ed. K. M. Carley. Pittsburgh, Pennsylvania: Carnegie Mellon University.

Gavetti, G. and D. A. Levinthal. 2000. Looking forward and looking backward: cognitive and experiential search. *Administrative Science Quarterly* 45: 113-137.

Hazy, J. K. and B. F. Tivnan. 2003. Representing Agent Intelligence as Local Network Dynamics. In *Proceedings of the 2003 North American Association of Computational Social and Organization Science*, ed. K. M. Carley. Pittsburgh, Pennsylvania: Carnegie Mellon University.

Hazy, J. K., B. F. Tivnan, and D. R. Schwandt. 2002. Boundary Spanning in Organizational Learning: Preliminary computational explorations. In *Proceedings of Managing the Complex IV Conference*, ed. M. Lissack. Fort Myers, Florida: Institute for the Study of Coherence and Emergence.

Hazy, J. K., B. F. Tivnan, and D. R. Schwandt. 2003. Permeable Boundaries in Organizational Learning: Preliminary computational modeling explorations. In *Proceedings of the Organizational Learning and Knowledge - 5th International Conference*, ed. M. Easterby-Smith. Lancaster, United Kingdom: University of Lancaster.

Hazy, J. K., B. F. Tivnan, and D. R. Schwandt. Under review. Modeling Social Structure as Network Effects: Computational Evidence That Rewarding Learning Improves Performance. Ashburn, Virginia: The George Washington University.

Krackhardt, D. and K. M. Carley. 1998. A PCANS Model Structure in Organization. Paper presented at the 1998 International Symposium on Command and Control Research and Technology, Monterey, California.

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

Sallach, D. L. 2003. The Simulation of Emergence: Mysticism, Reduction, or Social Complexity? In *Proceedings of the 2003 North American Association of Computational Social and Organization Science*, ed. K. M. Carley. Pittsburgh, Pennsylvania: Carnegie Mellon University.

Siggelkow, N. and J. W. Rivkin. 2003. Performance Determinants of Organizational Design: Towards an Understanding of Environmental Contingency. In *Proceedings of the 2003 North American Association of Computational Social and Organization Science*, ed. K. M. Carley. Pittsburgh, Pennsylvania: Carnegie Mellon University.

Tsvetovat, M. and K. M. Carley. 2003. Bouncing Back: Recovery Mechanisms of Covert Networks. In *Proceedings of the 2003 North American Association of Computational Social and Organization Science*, ed. K. M. Carley. Pittsburgh, Pennsylvania: Carnegie Mellon University.

AUTHOR BIOGRAPHIES

JAMES K. HAZY is a principal in a private capital advisory firm in New York and was formerly a financial VP at AT&T and CFO of an Ernst and Young, LLP business. Jim is a doctoral candidate in the Executive Leadership Program at the George Washington University. He earned an MBA with distinction from the Wharton School of University of Pennsylvania and a BS degree in mathematics from Haverford College. Jim's research interest is in the nature of emergent leadership activities in complex social systems. His email address is <Jim.Hazy@att.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 a 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 primary research interest addresses the co-evolutionary dynamics of strategic networks. His email address is <BTivnan@mitre.org>.