Social Networks and Multi-agent Organizational Performance

Matthew E. Gaston and **Marie desJardins**

Department of Computer Science
University of Maryland Baltimore County
{mgaston, mariedj}@cs.umbc.edu

Abstract

In studying the dynamical behavior of processes in artificial or natural social systems, a key factor is the topology of the network structure. It has been shown that real-world social networks tend to have non-random network structure with properties such as short average path length, excess clustering, and skewed degree distributions (Albert & Barabási 2002; Newman 2003). We show in this paper that the nature of the network structure in a social network of artificial or simulated agents has a significant effect on the performance of the overall system. We conclude that finding "good" network structures for a particular application domain is critical to modeling artificial social systems and implementing multi-agent systems. We argue that techniques for adapting network structure will be critical in large-scale agent communities.

Introduction

In recent years, the importance and ubiquitousness of networks has become increasingly apparent. The most obvious examples of this are the World-Wide Web (WWW) and the Internet. These networks have connected people, organizations, and companies in ways that have drastically affected international culture and the global economy. Several discoveries have been made about the underlying structure of the Internet and the World-Wide Web, including "diameter" approximations. Scientists have also proposed general principles, including "preferential attachment," for growing such networks (Albert & Barabási 2002).

Similarly, the underlying structure of social networks has been of interest to researchers for many years, starting with the work of Stanley Milgram (1967). Milgram was interested in investigating the lengths of paths between people in large social networks. Although his experiments were not comprehensive, his hypothesis that social networks have small diameter has been validated. In fact, many social networks have since been investigated and are found to support Milgram's low-diameter hypothesis, including scientific collaboration networks (Newman 2001) and telephone call graphs (Aiello, Chung, & Lu 2000).

Although much has been done to identify and exploit the salient structure of real-world networks, little has been done to examine the effects that these structures have on system

dynamics (Strogatz 2001; Watts & Strogatz 1998). The contribution of this paper is an empirical demonstration, in several different domains, that network structure has a significant effect on the dynamics of agent systems. The three systems examined are innovation diffusion in an artificial agent organization, opinion formation among networked artificial agents, and on-line multi-agent team formation. Our findings are relevant for both artificial societies (Carley & Gasser 1999; Epstein & Axtell 1996) and multi-agent systems (Huhns & Stephens 1999; Lesser 1999).

Background

Social networks, which have attracted the attention of scientists for many decades, are modeled and analyzed with the tools of graph theory. Three of the most prominently observed properties of real-world graphs are short average path lengths, excess clustering, and skewed degree distributions (Albert & Barabási 2002; Newman 2003; Strogatz 2001). Commonly referred to as the "small-world effect," *short average path length* was first recognized by analyzing the number of hops it took for postal mail to get to a specific destination if sent out randomly at first (Milgram 1967). Average path length is measured by calculating the average distance between any two nodes in the network (Watts & Strogatz 1998; Newman 2001).

Excess clustering is found in networks that have more clustering than would be expected in a random graph of the same size. The amount of clustering in a graph is defined as the number of three-cliques, or triangles, that exist in the graph normalized by the number of possible triangles (i.e., the number of connected triples of nodes) (Newman 2003).

The *degree distribution* of a network is the frequency of occurrence of nodes with each degree (i.e., number of connections). Although there are many possible distributions for degree, it has been observed that most real-world networks have a highly skewed degree distribution. These distribution have "heavy tails," and in general follow a power law, that is $P(k) \propto k^{-\gamma}$, which implies that the network has a hub and spoke structure, with some nodes having very large degree (Albert & Barabási 2002).

Regular graphs and random graphs are typically used to study artificial social systems. More recently, new network models have been developed in an attempt to capture properties observed in real-world networks. These models include

the small-world (Watts & Strogatz 1998) and scale-free (Albert & Barabási 2002) networks. All of these structures are briefly introduced in the following subsections.

Regular Graphs. The two-dimensional lattice, or grid, is commonly used to model the interaction topology of artificial social systems (Epstein & Axtell 1996; Schelling 1978; Latané 1981). In this interaction topology, agents have a geometric locality and are limited to directly interacting with agents that are located within a close proximity. The two-dimensional lattice does not exhibit short average path lengths or a skewed degree distribution, but does exhibit a form of generalized clustering, called transitivity. Transitivity is measured by considering the number of shared neighbors among nodes (Newman 2003). As will be seen below, lattice networks are the foundation of small-world networks.

Random Graphs. Random graphs were first introduced in the 1960s by Erdös and Rényi (Erdos & Renyi 1959). A *random graph* $G_{N,p}$ consists of N nodes that are connected randomly, with p denoting the probability of an edge existing between a randomly chosen pair of vertices. Random graph models have been widely studied and are quite useful since many of their properties can be computed analytically. For instance, the average number of undirected edges in $G_{N,p}$ is $N(N-1)p/2$, and the average degree of a vertex is $k = p(N-1) \approx pN$, where the approximate solution holds for large N.

Recent evidence suggests that random graphs are not representative of real-world networks (Albert & Barabási 2002). They do possess short average path lengths, but random graphs do not exhibit large amounts of clustering or a skewed degree distribution. Random graphs are included in this study, because they have been widely used for decades. We will compare the dynamics that arise from random graphs to those arising from alternative graph models that are structured more like real-world networks.

Small-world Networks. The small-world network model is an attempt to introduce more clustering into networks and to account for observed short average path lengths (Watts & Strogatz 1998). A key observation is that small-world networks have properties that lie between those of regular (lattice) networks and random graphs.

Lattices are the starting place for small-world networks. Lattice models lend themselves directly to increased clustering because they imply physical localities among agents. A simple example of this is a ring of nodes (i.e., a one-dimensional lattice wrapped upon itself) where each node connects to the closest nodes on either side. The key feature here is that for any neighborhood, most of the nodes are connected to one another inducing clustering.

The mechanism for decreasing average path length is a probabilistic randomization of edges, resulting in shortcut connections through the graph as seen in Figure 1. The parameter ρ is used to determine if an edge is replaced by a random shortcut through the graph. One caveat in this construction of a small-world graph is that the graph can be-

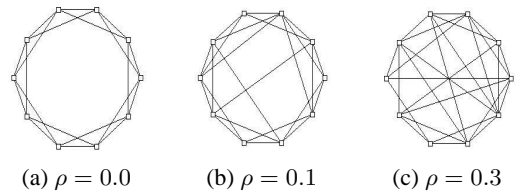(a) $\rho = 0.0$ (b) $\rho = 0.1$ (c) $\rho = 0.3$

Figure 1: The graphic shows three increasingly random small-world networks: (a) shows a one-dimensional lattice with no shortcut links; (b) shows the same lattice with a few shortcuts; and (c) shows a small world with many shortcuts, which begins to resemble a random graph. All three of the networks are constructed from a one-dimensional lattice where nodes are connected to four other nodes based on physical proximity.

come disconnected. When edges are replaced by shortcuts with probability $\rho = 1$, the graph is a random graph.

Scale-free Graphs. The scale-free graph model for networks is motivated by the degree distributions of the Internet and the WWW (Albert & Barabási 2002). The model is a highly intuitive model based on how networks are believed to evolve and grow in the real world.

The generation of scale-free graphs has two simple rules:

1. **growth:** at each time step, a new node is added to the graph, and

2. **preferential attachment:** when a new node is added to the graph, it attaches preferentially to nodes with high degree.

Preferential attachment is modeled by

$$P(e_{ij}) = mk_j \sum_{v \in V} (\frac{1}{k_v})^\alpha \qquad (1)$$

where i is a new node being added to the network and $P(e_{ij})$ is the probability of the creation of an edge from i to j. Here, k_v is the degree of node v. There must be an initial connected core of m_0 nodes, ensuring that at the beginning of the graph generation process, nodes with non-zero degree exist. There are several additional parameters for the model including the number of expected initial connections that a node will make m and a scaling factor for the probability of connection α to force the graphs to be more or less dense. Note that when $m = \alpha = 1$, the probabilities of connecting to existing nodes in the network sum to exactly one. Scale-free graphs exhibit both short average path lengths and a skewed degree distribution but, like random graphs, lack excess clustering (Albert & Barabási 2002).

Models and Results

We explored the effect of network structure on the dynamics of three different models of social systems and multi-agent systems: innovation diffusion, opinion formation, and team formation. These models were chosen because they are widely studied in the computational organization theory and

multi-agent systems literature. In each of the experiments described below, the performance of the system was measured for four different network structures of equal density: a two-dimensional lattice, a two-dimensional small-world network (connectedness verified), a random graph (connectedness verified), and a scale-free graph. In all of the experiments, the lattice network was used to determine the size and the density of connections in the agent organization and the other network models were appropriately parametrized to ensure the same size and density. Due to space limitations, we have presented only representative results, but the general findings hold across a wide range of networks.

Innovation Diffusion

Innovation diffusion, taken in the context of the *theory of the firm*, is the process by which an organization adopts new ideas. Organizational structure plays a key role in an organization's innovation diffusion capability. DeCanio and Watkins (1998) introduced a model of innovation diffusion that depends heavily on the processing capability of the agents in the organization.

In the initial study, the effects of structure on an organization's ability to diffuse innovation were presented. The network structures studied in the original work included a fully connected network, a designed hierarchy, a hypercube (n-dimensional lattice), and a small random graph. We extended the study to include much larger networks that have more realistic structural properties.

Model. The basic model places n agents into an organizational structure where the connections among agents determine who influences whom. The agents take on binary state values $\{0, 1\}$ representing whether or not a given agent has adopted the innovation. Initially, none of the agents have adopted the innovation. As the system evolves, agents update their state based on the ratio of their neighbors that have adopted the innovation. One agent is selected to initiate the innovation. Once an agent adopts the innovation, they do not change state again.

The probability that an agent i changes state is:

$$P_i = f(x_i), \quad \text{where} \quad x_i = \sum_{j \in V} \frac{e_{ij} y_j}{k_i}. \quad (2)$$

In (2), $y_j = \{0, 1\}$ is the state of agent j, k_i is the degree of agent i, and $e_{ij} = 1$ if i is connected to j and 0 otherwise. Agent i then updates its state by

$$y_i = \begin{cases} 1 & \text{with probability} \quad P_i \\ 0 & \text{with probability} \quad 1 - P_i. \end{cases} \quad (3)$$

Time is discrete and the agents update synchronously.

The function f is a function of the agents' processing capability and of the ratio of the agent's neighbors that have adopted the innovation. Information processing capability can be modeled using a modified logistic function (DeCanio & Watkins 1998).[1] An agent's readiness to switch states

---

[1] Other *squashing* functions would apply; for consistency with the work of DeCanio, the modified logistic function is used.

(i.e., to adopt the innovation) is based on the scaled ratio of its neighbors' states and the agent's own processing capability. The modified logistic function is

$$f(x) = 1/[1 + e^{[x-a/b]/b/c}] - 1/1 + e^{a/b}. \quad (4)$$

The processing capability of an individual agent effectively models the agent's cognitive capacity. The lower an agent's processing capability, the harder it is for the agent to process information from its neighbors. Given perfect information processing capabilities ($c = 1.0$) the fully connected network will always perform the best (i.e., all information is avaible to all agents at all times and can be processed perfectly). When processing capabilities drop below perfect, the number of connections among the agents make it more difficult to process all of the available information.

Results. In the original study, all of the agents in the system have the same processing capability (DeCanio & Watkins 1998). We maintain that assumption for the results presented here, but the general findings hold true for agents with heterogeneous processing capabilities. In these experiments, the number of agents is 10,000 and each data point is the average of 25 simulations.
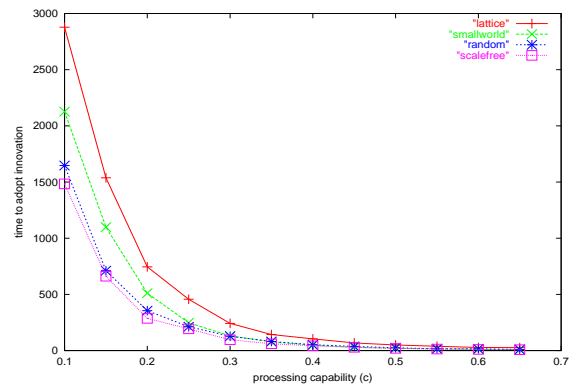
Figure 2: Time for the entire system to adopt an innovation as a function of processing capability.

In support of the original results (DeCanio & Watkins 1998), the most hierarchical network structure, the scale-free graph, performs the best (i.e., adopts the innovation the fastest) for all processing capabilities as seen in Figure 2. Next is the random graph, followed by the small-world network. The nature of the innovation diffusion model is that it resides in communication, or information, space. One of the factors contributing to the somewhat poor performance of the small-world model may be its ties to geographic space. The result show that the scale-free graph structure supports rapid diffusion of innovation in a large organization of agents with sub-optimal information processing capabilities.

Opinion Formation

Opinion formation is a popular topic in the social science community. Not unlike innovation diffusion, opinion forma-

tion models address the question of how opinions (usually binary, but sometimes continuous) can spread and change through an agent organization. There are many different factors that affect organizational opinion formation. One factor is that of *social impact* based on a notion of capturing two agents' *social distance*: the distance between the two agents' beliefs and ideals (Latané 1981).

Model. The social impact model of opinion formation and reformation is loosely based on the Ising model from condensed matter physics (Latané 1981). In the this model, there are n agents. Each agent, i, possesses a strength, $s_i > 0$, an opinion, $\sigma_i \in \{-1, 1\}$, and a position in some defined space. In the original model, the position of each agent was fixed in a two-dimensional grid (Latané 1981). In this situation, agent i's social distance from agent j is simply the Euclidean distance between their respective locations. Also included in the model is a notion of an external influence, h, which can be thought of as information external to the organizations that the agents can observe. Similarly, social noise, η, is included in the model and serves as a mechanism for including uncertainty, trust, and reputation in the dynamics of the model.

Social impact is the impact of the entire organization on an agent's opinion formation process (Latané 1981). To calculate social impact, the partial impact for agent i is first calculated by

$$I_i = \sum_{j=\{1,\dots,n\}} \frac{s_j \sigma_j}{d_{ij}} + h, \qquad (5)$$

where $d_{ii} = 1$. The equation for partial impact captures the weighted sum of the impact of all of the agents relative to the location of agent i with h being external information. Once the partial impact is calculated, it is mapped to a probability that accounts for social noise, by

$$p(I_i) = \frac{e^{I_i/\eta}}{e^{I_i/\eta} + e^{-I_i/\eta}}. \qquad (6)$$

This function is a squashing function to create non-linear effects on changing opinions. The agent updates its opinion based on this probability where

$$\sigma_i = \begin{cases} +1 & \text{with probability} \quad p(I_i), \\ -1 & \text{with probability} \quad 1 - p(I_i). \end{cases} \qquad (7)$$

The social impact opinion formation model does not account for the social network structure among the agents explicitly. The agents in the original model are in fact fully connected; every agent impacts every other agent's opinion. To explore the effects of network structure on opinion formation, we replace the use of the Euclidean distance metric with that of geodesic (shortest path) distance measured from the agent social network structure. The results presented below are provided in the context of a specified leader attempting to change the opinion of the entire organization.

Results. A *leader*, in the context of opinion formation, is an agent that can drive the organization toward a given opinion (i.e., an agent that can make every other agent agree with

their opinion). Leaders must be strong to convince an entire organization to change its opinion. The underlying structure of the organization impacts the necessary strength of the leader.

As mentioned above, we replace the Euclidean distance with the normalized length of the shortest path between two agents for determining social distance. The lattice imposes a distance measure similar to physical distance, although it results in a Manhattan distance measure rather than an Euclidean distance measure.

Using the network to define social distance, simulations were conducted to determine the strength required of the leader to force all of the other agents to change their opinion. The agents are given random initial opinions. After 10 simulation time steps, the designated leader switches and fixes its opinion. Once 90% of the agents in the organization change to the leader's opinion, the time is recorded. If 90% of the organization does not change opinion after 100 time steps, it is determined that the leader is not strong enough to change the opinion of the entire organization. In all simulations the leader is the most highly connected agent (i.e., the agent with the smallest average distance to all other agents). Based on the original studies (Latané 1981), the number of agents is 625; social noise $\eta = 167.3$; external influence $h = -0.5$; and each data point is the average of ten simulations.
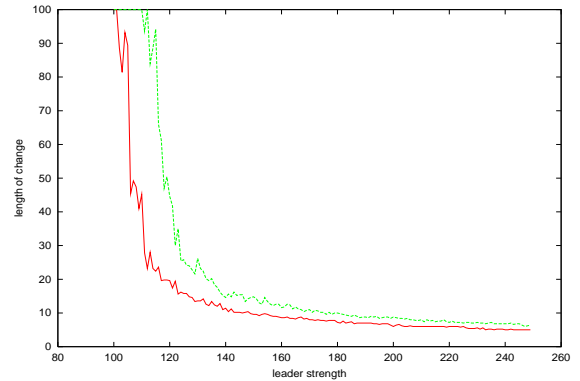

Figure 3: Time to shift the entire network's opinion, as a function of leader strength, for scale-free networks (solid red) and random graphs (dashed green).

Neither the lattice nor the small-world network ever reached 90% of the agents agreeing with the leader. In both of these network structures, agents exist that are far enough away from the leader that the social noise diffuses the influence of the leader's opinion.

The results for using the random graph and the scale-free graph to define social distance are given in Figure 3. The random graph requires a stronger leader than that of the scale-free graph. This is explained by the fact that in the scale-free graph, the leader has shorter average social distance to all other agents.

Team Formation

To explore the effects of network structures on team formation, we distill a simple agent-based organizational model

from previous work on multi-agent team formation (Nair, Tambe, & Marsella 2002; Abdallah & Lesser 2004). Tasks are generated and globally advertised to the agents in the organization, and agents form teams to accomplish the tasks. The network structure restricts which agents can be on teams together. The tasks are generic in that they only require that a team of agents with the necessary skills form to accomplish the specific task.

The organization contains a set of N agents situated in a specified social network. Each agent can be in one of three states: uncommitted, committed, or active. An *uncommitted* agent is available and not assigned to any task. A *committed* agent has chosen a task, but the full team to work on the task has not yet formed. Finally, an *active* agent is a member of a team that has fulfilled all of the skill requirements for a task and is actively working on that task. The agents are also assigned a single skill, $\sigma \in [1, \Sigma]$, where Σ is the number of different types of skills that are present in the organization.

Tasks are introduced at fixed task introduction intervals, where the length of the interval between tasks is given by a parameter, μ. Tasks are globally advertised (i.e., announced to all agents). Each task T has an associated size requirement, $|T|$, and a $|T|$-dimensional vector of required skills, R_T. The skills required for a given task T are chosen uniformly from $[1, \Sigma]$. Each task is advertised for a finite number of time steps proportional to its size (namely $\delta|T|$, where δ is a model parameter) to ensure that the resources (i.e., agents) committed to the tasks are freed if the full requirements of the task cannot be met. Similarly, teams that form to fill the requirements of a given task are only active for a finite number of time steps (namely $\alpha|T|$, where α is a model parameter).

Definition: A **valid team** is a set of agents $\{a_i\}$ whose corresponding set of nodes $\{v_i\}$ induce a *connected subgraph* of G and whose skill set $\{\sigma_i\}$ fulfills the skill requirements for a given task T.

The restriction of a team forming a connected subgraph of the organizational structure forces the agents to join teams based on local knowledge using heuristics for the likelihood of a given partial team's success. The local knowledge available to each agent includes the number of positions filled on the team, the number of uncommitted immediate neighbors of the agent, and the number of immediate neighbors on a specific team. In the experiments described here, when deciding what team to join, agents joined teams based on the definition of a valid team (i.e., one adjacent agent must be on the team), with a probability proportional to the number of filled positions on the team (i.e., agents prefer to join teams that are nearly completely filled). Agents create teams for new tasks probabilistically based on the number of uncommitted immediate neighbors.

Results. We conducted a series of simulated experiments to measure the effects of the agent interaction topology on the dynamics of multi-agent team formation. Here, we focus on the task introduction interval μ. In these experiments,

the number of agents is 100; the simulation length is 5000 iterations; and each data point represents the average of 100 simulations.

The measure that we use in our experiments for global performance is the *organizational efficiency* of the agent society, defined as:

$$\text{organizational efficiency} = \frac{\text{\# of teams successfully formed}}{\text{total \# of tasks introduced}}.$$

(8)

This measure of organizational efficiency captures the overall performance of the networked organization of agents, as we are directly interested in the cooperative distributed performance. Because tasks are introduced periodically at a deterministic rate, this measure is fair across all network structures and experiments.
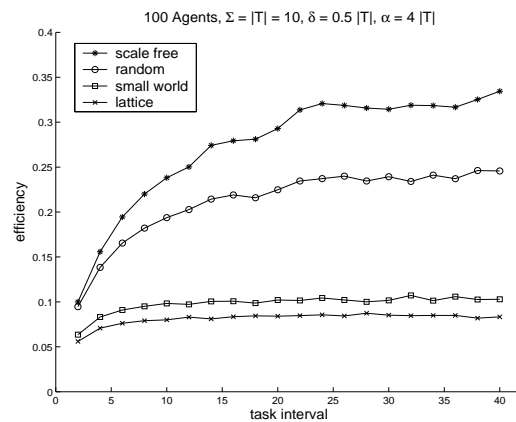

Figure 4: Organizational efficiency vs. task interval.

Figure 4 shows the organizational efficiency of the agent societies as a function of the task introduction interval μ for each of the four network structures: scale-free, random, small-world, and lattice. Agents connected with a scale-free network structure ($-\diamond-$) significantly outperformed the other three network structures.[2]

Although the scale-free network dominates, there are several important performance differences among the other structures. First, in general, the more stochastic network structures (i.e., the scale-free and the random networks) consistently outperform the more regular network structures (i.e., the lattice and the small-world). This pattern is replicated in additional experiments. Second, within the "regular" graphs, the small-world network outperforms the lattice network structure, despite their high degree of similarity. In these experiments, the small-world network was constructed by simply rewiring edges of the lattice network with a fairly low probability, $p = 0.05$. Therefore, 95% of the edges in these two network structures are identical. This 5% rewiring led to a statistically significant improvement in organizational efficiency.

---

[2]These results are statistically significant at the 99th percentile confidence level, using a student t-test.

Related Work

In the artificial intelligence community, a small selection of studies have emphasized the importance of social structure on organizational behavior. The spread of social conventions in a society of agents is affected by the complex network structure among the agents (Delgado 2002). Line-of-sight considerations affect the performance of mobile robot teams (Arkin & Diaz 2002). In the study of social system dynamics and agent-based modeling, the underlying network structure of the system has been shown to significantly impact the formation of firms (Axtell 2000). More recently, research on multi-agent games in complex network structures demonstrated that variations in agent social networks lends itself to significantly different stabilities in agent strategies (Abramson & Kuperman 2001; Holme *et al.* 2003).

Conclusions and Future Work

We have demonstrated that the social network structure underlying an agent organization can have significant impact on organizational performance. Given this result, future work should be devoted to identifying efficient network structures and, in particular, decentralized organizational learning techniques based on local network adaptation.

Preliminary investigations and findings include network adaptation for distributed information retrieval (Yu, Venkatraman, & Singh 2003), peer-to-peer networking (Ramanathan, Kalogeraki, & Pruyne 2002), and dynamic multi-agent team formation (Gaston & desJardins 2004). In addition, other agent-level learning techniques may be developed that allow agents embedded in a social network to learn from experience an improve organizational performance despite the network structure. For now, while traditional artificial intelligence research focuses on what an agent knows, we have shown that who an agent knows is also important.

References

Abdallah, S., and Lesser, V. 2004. Organization-based cooperative coalition formation. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Techonology (IAT)*.

Abramson, G., and Kuperman, M. 2001. Social games in a social network. *Phys. Rev. E* 63(030901).

Aiello, W.; Chung, F.; and Lu, L. 2000. A random graph model for massive graphs. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, 171–180.

Albert, R., and Barabási, A. 2002. Statistical mechanics of complex networks. *Review of Modern Physics* 99(3):7314–7316.

Arkin, R. C., and Diaz, J. 2002. Line-of-sight constrained exploration for reactive multiagent robotic teams. In *Proceedings of the Seventh International Workshop on Advanced Motion Control (AMC'02)*.

Axtell, R. 2000. Effects of interaction topology and activation regime in several multi-agent systems. In *MABS*, 33–48.

Carley, K., and Gasser, L. 1999. Computational organization theory. In Weiss, G., ed., *Introduction to Multiagent Systems*. Morgan Kaufmann. Chapter 7.

DeCanio, S., and Watkins, W. 1998. Information processing and organizational structure. *Journal of Economic Behavior and Organization* 36:275–294.

Delgado, J. 2002. Emergence of social conventions in complex networks. *Artificial Intelligence* 141:171–185.

Epstein, J., and Axtell, R. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. The Brookings Institute Press and MIT Press.

Erdos, P., and Renyi, A. 1959. On random graphs. *Publicationes Mathematicae* 6:290–297.

Gaston, M., and desJardins, M. 2004. Adapting network structure for efficient team formation. In *Proceedings of AAAI 2004 Fall Symposium on Artificial Multi-agent Learning*.

Holme, P.; Trusina, A.; Kim, B. J.; and Minnhagen, P. 2003. Prisoners' dilemma in real-world acquaintance networks: Spikes and quasi-equilibbria induced by the interplay between structure and dynamics. *Phys. Rev. E* 68(030901).

Huhns, M., and Stephens, L. 1999. Multi-agent systems and societies of agents. In Weiss, G., ed., *Multi-Agent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press.

Latané, B. 1981. The psychology of social impact. *American Psychol.* 36.

Lesser, V. 1999. Cooperative multi-agent systems: A personal perspective on the state of the art. *IEEE Trans. Knowledge and Data Engineering*. Jan.-Feb. Special Issue.

Milgram, S. 1967. The small world problem. *Psychology Today* 22:61–67.

Nair, R.; Tambe, M.; and Marsella, S. 2002. Team formation for reformation. In *Proceedings of the AAAI Spring Symposium on Intelligent Distributed and Embedded Systems*.

Newman, M. 2001. The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci.* 98:404–409.

Newman, M. 2003. The structure and function of networks. *SIAM Review* 45:167–256.

Ramanathan, M. K.; Kalogeraki, V.; and Pruyne, J. 2002. Finging good peers in peer-to-peer networks. In *International Parallel and Distributed Computing Symposium (IPDPS)*.

Schelling, T. C. 1978. *Macromotives and Microbehavior*. WW Norton and Co.

Strogatz, S. 2001. Exploring complex networks. *Nature* 410:268–276.

Watts, D., and Strogatz, S. 1998. Collective dynamics of 'small-world' networks. *Nature* 393:440–442.

Yu, B.; Venkatraman, M.; and Singh, M. P. 2003. An adaptive social network for information access: Architecture and experimental results. *Applied Artificial Intelligence* 17(1):21–38.