

A Simple Learning Approach for Endogenous Network Formation

Matthew E. Gaston and Marie desJardins*

Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
{mgastol, mariedj}@cs.umbc.edu

Abstract

We operationalize the theory associated with an economic model of network formation and conduct experiments to determine the feasibility of endogenous, dynamic network formation in multi-agent organizations. We develop a learning-based, decentralized network formation strategy that allows the agents to make network adaptation decisions based on past performance. We compare our method with a dynamic network formation process proposed in the economics literature that relies on a global computation over the entire network structure. Our findings demonstrate that local decisions based solely on prior experience perform as well as local decisions based on perfect knowledge and a global computation.

Introduction

As multi-agent systems grow in size and complexity it is likely that they will move toward an open system paradigm, within which neither the system nor all of the individual agents will fall under the control of a single authority. There are many challenging problems associated with open multi-agent systems, one of which is the ability of the agents to manage their local connectivity in a large agent social network. In many multi-agent systems, it is assumed that all agents can interact with all other agents and that they can know about the other agents and their capabilities. In the open system paradigm, especially when there is a large number of agents, these assumptions will no longer hold. Limited cognitive capabilities, geographical constraints, and communications limitations are three factors that will contribute to restricted interactions among the agents and the emergence of an agent social network in large, open multi-agent societies. Examples of networked multi-agent system domains include distributed information retrieval (Yu, Venkatraman, & Singh 2003), supply chains (Thadakamalla *et al.* 2004; Walsh & Wellman 2003), sensor networks (Culler, Estrin, & Srivastava 2004), and team formation (Nair, Tambe, & Marsella 2002).

One potential method of collective learning in a large network of agents is to endow the agents with the ability to

adapt, or learn to adapt, their local connectivity structure in the agent social network. We refer to the concept of distributed, local, real-time adaptation of the agent social network as *agent-organized networks* (AONs). This notion of organizational learning as a result of AONs is believed to be widely applicable to many multi-agent system applications, and in particular, systems that have one or more of the following characteristics:

- it is difficult for agents to know about or interact with all other agents in the system,
- agents can enter and leave the multi-agent society,
- no single authority controls the entire system,
- no single authority controls all of the agents,
- there is a cost (either explicit or implicit) associated with agent interactions, and
- interaction among the agents leads to changes in performance or value (for the agents, the network, or both).

This paper does not attempt to fully develop the notion of AONs, but demonstrates the need and importance of AONs as a distributed learning mechanism.

In this paper, we review the economic theory of a network formation game and use this theory to discuss the challenges of forming networks based on the decisions of individual agents. The main contributions of this paper include the identification of distributed network adaptation as a mechanism for cooperative learning in networked multi-agent systems and the development and exploration of a local, distributed learning mechanism that allows agents to effectively form social networks. We conclude the paper by relating our work to previous work in multi-agent systems and by suggesting several promising future directions for multi-agent network formation and adaptation.

The (Symmetric) Connections Model

In this section we introduce an economic network formation game, *The Symmetric Connections Model*, and briefly survey the notions of stability and efficiency for the model.

The Model

First presented by Jackson and Wolinsky (1996), *The Connections Model* (CM) is a stylized model that is representative of a larger class of network games. In such games,

*The authors thank the anonymous reviewers and Blazej Bulka for helpful comments and insights.
Copyright © 2005, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

values are given to network structures and to the positions of agents in the network. CM was first developed in order to characterize and study the nature of “social communications among individuals” (Jackson & Wolinsky 1996).

In the model, agents directly communicate with the agents with whom they share an undirected edge (i.e., a connection) in the network structure. The communications imply value for information flow, so agents also benefit from the indirect communications represented by their neighbor’s direct connections and their neighbor’s neighbor’s connections. The value of the indirect communications falls off as a function of the geodesic distance (i.e., shortest path distance) in the network. In particular, the value allocated to agent i in network G in CM is given as

$$Y_i(G) = \sum_{j \neq i} \delta_{ij}^{d(i,j)} - \sum_{j:ij \in G} c_{ij}, \quad (1)$$

where $ij \in G$ denotes a connection between agents i and j in the network G , and $d(i, j)$ is the shortest path distance between i and j . The parameters δ_{ij} , with $0 < \delta_{ij} < 1$, are the values of the (possibly indirect) connections between agents i and j discounted as a function of the distance between the two agents. The parameters c_{ij} are the costs of direct connections between agents i and j . Agents benefit from being “close” to other agents, discounted by distance, while they only suffer costs for their direct connections. The value of a network in the CM is

$$v(G) = \sum_i Y_i(G). \quad (2)$$

The model allows agents to create or remove connections, with the goal of maximizing $Y_i(G)$. The model assumes that connections must be added bilaterally (i.e., the agents at both ends of the connection must agree to the connection), but that connections can be removed unilaterally. The *Symmetric Connections Model* (SCM) is symmetric in that it has homogeneous values for all of the connections: $\forall ij \in G$ $\delta_{ij} = \delta$ and $c_{ij} = c$.

CM and SCM are representative of a larger set of network games that have been considered in the economics literature. Network formation has also been studied in the context of trade networks, labor markets, coauthor networks, and buyer-seller networks (Jackson 2003). Much of the economics literature is concerned primarily with stability, efficiency, and equilibrium, and does not concern itself with the real-time, dynamic behavior of agents in such models. Before considering the dynamic, real-time behavior of agents in the SCM, we first present the primary stability and efficiency results for the SCM.

Stability and Efficiency

As mentioned above, the economic literature is primarily concerned with stable and efficient structures in network formation games. In this section, we present the major theoretical results on stability and efficiency in the SCM.

Stability. The notion of stability captures whether or not agents in the network desire to make any single change to the network structure.

Definition 1 (Jackson & Wolinsky 1996)¹ A network G is **pairwise stable** with respect to allocation rule Y if

- (i) $\forall ij \in G$, $Y_i(G) \geq Y_i(G - ij)$ and $Y_j(G) \geq Y_j(G - ij)$, and
- (ii) $\forall ij \in G$, if $Y_i(G + ij) > Y_i(G)$ then $Y_j(G) > Y_j(G + ij)$.

In the definition, $G - ij$ and $G + ij$ represent the removal of the connection between agents i and j and its addition, respectively.

Intuitively, pairwise stability implies that no agent desires to make any *one* modification to the connections in the network (i.e., no agent desires to add or delete a connection). The emphasis on a single modification, known in the game theory literature as *myopic*, is important as agent strategies that rely on the notion of pairwise stability can result in locally optimal network structures (see below).

Efficiency. A pairwise stable network is a network for which no agent desires to change any one of its connection. Efficiency is a more strict notion.

Definition 2 (Jackson & Wolinsky 1996) A network G is **efficient** if $v(G) \geq v(G') \forall G' \in \mathcal{G}$.

Here, \mathcal{G} is the space of all network structures of a particular size (i.e., the number of agents in the system). In essence, an efficient network has a value that is at least as great as *any other* network structure.

Now that we have introduced the notions of stability and efficiency, we present the two major network formation results from Jackson and Wolinsky (1996).

We start with a proposition regarding the pairwise stability of networks under different parameter regimes.

Proposition 1 (Jackson & Wolinsky 1996) *In the symmetric connections model:*

- (i) A pairwise stable network has at most one (non-empty) component.
- (ii) For $c < \delta - \delta^2$, the unique pairwise stable network is the complete graph, C^N .
- (iii) For $\delta - \delta^2 < c < \delta$, a star encompassing all players is pairwise stable, but not necessarily the unique pairwise stable network.
- (iv) For $\delta < c$, any pairwise stable network which is nonempty is such that each player has at least two links and thus is efficient.

Here, a star network is one with every agent connected to a central hub, and no other connections. The proof of the proposition is given in the original study (Jackson & Wolinsky 1996), but as our study in later sections focuses on case (iii) of the proposition, we provide the intuition behind its proof. When $\delta - \delta^2 < c < \delta$, the star is pairwise stable because the hub would not delete any of its connections

¹This definition is taken directly from Jackson and Wolinsky (1996) with a slight change of notation due to the assumption that Equation (2) gives the value function of the network.

as $\delta > c$, and none of the other agents would form a direct connection because for any two of the agents i and j : $d(i, j) = 2$, agent i gets the δ^2 indirect benefit via its connection with the hub, and $\delta^2 > \delta - c$ (where $\delta - c$ is the net cost of the new direct connection between i and j).

We have seen the result for the structure of pairwise stable networks, now we present the result for the structure of efficient network structures.

Proposition 2 (Jackson & Wolinsky 1996) *The unique efficient network structure in the symmetric connections model is*

- (i) *the complete graph, C^N , if $c < \delta - \delta^2$,*
- (ii) *a star encompassing everyone if $\delta - \delta^2 < c < \delta + \frac{(N-2)}{2}\delta^2$, and*
- (iii) *no links if $\delta + \frac{(N-2)}{2}\delta^2 < c$.*

The proof of the proposition can be found in the original paper (Jackson & Wolinsky 1996) and we provide the intuition for case (ii). The star is the unique efficient structure for two reasons: 1) it has the minimum number of connections, $(n - 1)$, to guarantee a single component, therefore minimizing cost, and 2) those connections are arranged so as to minimize the average pairwise distance between all of the agents (i.e., the star minimizes the mean path length, or diameter, of the network), maximizing the benefit to all of the agents.

The theory described here is extremely useful in understanding the structure of networks and network formation. Although simple, the SCM presents a challenging problem for distributed, multi-agent cooperations. Local decisions that increase local utility can decrease the utilities of other agents in the network and in turn decrease the overall value of the network. Our main concern is in designing strategies (either directly or learned) that allow agents to make local decisions in real-time as the network evolves in order to form efficient network structures.

On the Challenge of Distributed, Dynamic Network Formation

The theory of the SCM presented in the last section is useful when there is central control of the network structure. Our main concern in this paper is the operationalization of the SCM in a multi-agent environment where the agents make decisions locally in order to attempt to form efficient network structures. From here forward, we will assume that agents are cooperative, in that they do not make decisions to intentionally degrade the value of the network.

A Dynamic Network Formation Process

Using the notion of pairwise stability, Watts (2001) proposed a dynamic model for the network formation process (Jackson & Watts 2002). The dynamic model starts with an empty network (i.e., no connections), and then at each iteration, two agents, i and j , are chosen randomly from a probability distribution $p(i)$ and allowed to consider their connections (or lack thereof). We call this mechanism for considering connections the *random meeting* mechanism. For the

remainder of this paper, we will assume that $p(i)$ is the uniform distribution over all of the agents in the network.

The deterministic dynamic network formation process (Watts 2001) is as follows. Let G represent the graph before i and j consider their connection. If $ij \in G$, the agents remove the connection if $Y_i(G - ij) > Y_i(G)$ or $Y_j(G - ij) > Y_j(G)$. If $ij \notin G$, the agents add the connection if $Y_i(G + ij) \geq Y_i(G)$ and $Y_j(G + ij) \geq Y_j(G)$, with the inequality holding strictly for at least one of i or j . That is, the agents establish the connections if it is mutually beneficial, and they remove the connection if either benefits.

In order to prevent the model from remaining in “uninteresting” stable states (e.g., the network with no connections when $c > \delta$), a stochastic dynamic network formation process was also considered. The stochastic model is the same as above, although decisions to establish connections are randomly inverted with probability ϵ . This was meant to represent slight irrationality or small “tremors” in the deterministic process (Watts 2001).

Note that the Watts formation process requires an individual agent making a local decision to perform a global computation. That is, the agent computes the change in value (which depends on the entire network structure) after adding or deleting a connection, and then determines if the change to the connection should remain or be reversed based on whether the value increased or decreased, respectively. This assumption, that agents have global knowledge of the entire network structure for informing decisions, is not realistic in many multi-agent environments. At the same time, we will use the Watts formation process to emphasize the challenge of forming networks in a completely decentralized manner before presenting our new method for the decentralized formation of networks with only local information and past performance.

Applying the Watts Formation Process

When the dynamic network formation process was first proposed, it was realized that forming an efficient and pairwise stable network (i.e., the star when $\delta - \delta^2 < c < \delta$) was difficult when the number of agents grows large.

Proposition 3 (Watts 2001) *Consider the symmetric connections model in the case where $\delta - \delta^2 < c < \delta$. As the number of players grows, the probability that a stable state (under the process where each link has an equal probability of being identified) is reached with the efficient network structure of a star goes to 0.*

The proof of the proposition is based on the fact that no agent wants to bear the burden of being the hub node in the star network. The hub node in a star network has a much lower value than the other nodes in the network.

Given the finding that the unique, efficient, and pairwise stable network structure is very unlikely to be discovered, the question remains as to what network structures are found by dynamic network formation processes. We ran computational experiments using the Watts dynamic network formation process to study the structure of endogenously formed networks. A representative sample experiment is depicted in Figure 1.




Figure 1: The result of distributed network formation when the agents employ the global computation strategy when $\delta = 0.9$ and $c = 0.8$: a) the value of the network over time, b) various structural properties of the network as it evolves, and c) the resulting network topology after 50,000 iterations. The properties in b) are: *density* or ratio of the number of connections to all possible connections; *diameter* or average shortest path length between all pairs of nodes; *clustering* the ratio of triangles to connected triples; and *rewirings* or the cumulative number of connection changes.

Figure 1(a) shows the value of the network, $v(G)$, over time as the agents use the stochastic dynamic network formation process of Watts (2001), which we call the *global computation* strategy. In the experiment presented, there are 100 agents, $\delta = 0.9$, $c = 0.8$, and ϵ was initialized to 0.05 and slowly decreased over time. We experimented with a range of parameters in $\delta - \delta^2 < c < \delta$ and found similar results. The figure shows that the agents rapidly form a network and then the value of the network plateaus just under 7000. The optimal value for the star network with $n = 100$, $\delta = 0.9$, and $c = 0.8$ is 7878.42.

Clearly, the global computation strategy does not guarantee finding the optimal star network, as there is simply too much competition among the agents for one of them to become the hub (i.e., all of the agents are attempting to locally maximize $Y_i(G)$). At the same time, the agents are able to find a network structure that supports an even distribution of value and a large network value.² Of course, this is expected as the agents are making perfect decisions ($1 - \epsilon$ percent of the time) with perfect, and global, information about the network structure.

Forming Networks without a Global Computation

In this section, we propose an individual agent learning strategy that only requires local information and past performance in order to allow the agents to form networks in a decentralized fashion. In the design of our learning strategy, we wanted to uphold several assumptions that are likely to apply in many multi-agent environments:

- **no global knowledge:** our strategy does not use any global knowledge of the network structure or the other agents in the network,

²The sizes of the nodes in the network structure shown in Figure 1(c) are proportional to their value, $Y_i(G)$. There is little deviation in the size of the nodes in the network.

- **no explicit modeling of other agents:** our strategy does not require large amounts of memory for modeling the other agents in the environment, or values of any specific connection in the network,
- **synchronous:** following Watts (2001), we employ the random meeting mechanism, and for now, only one network adaptation can occur at any moment in time,
- **bilateral connection establishment:** both agents must agree to establish a connection,
- **unilateral connection removal:** a single agent can decide to remove any of its connections.

Our learning strategy is derived from stateless (i.e., single state) multi-agent Q -learning, as it attempts to model the utility of taking any action over time.

A Simple Learning Strategy

At any moment, the agents in the SCM can choose one of three actions: add a connection, delete a connection, or do nothing. The choice of the action is dependent on the agent that is met at random (either an existing connection or a new connection) and we extend the choice of the action to be dependent on the expected change in value for taking a certain action. As this is the first known attempt to apply online learning to endogenous, multi-agent network formation, we have selected a simple learning approach based on stateless Q -learning (Claus & Boutilier 1998).

Let $A = \{add, delete, nothing\}$ be the set of actions available to an agent. We update the action value function Q for agent i , given that agent i took action a , using

$$Q_i(a) \leftarrow (1 - \alpha)Q_i(a) + \alpha(\Delta Y_i(G)), \quad (3)$$

where α is the learning rate and $\Delta Y_i(G)$ is the change in an agents value after the action is taken. Note that this is in contrast to first computing the change and then taking the appropriate action. Our method makes the change, and then the value of the change is made available as a reinforcement signal. We update the Q values whenever an action

is taken. Additionally, we update $Q_i(\text{nothing})$ whenever $|\Delta Y_i(G)| > 0$ when agent i is idle (i.e., other agents in the network are changing connections that affect the value of agent i). This update to $Q_i(\text{nothing})$ is what makes our method a multi-agent learning approach. In essence, agents learn when changes elsewhere in the network are influencing their value (either positively or negatively). In all of the experiments presented in this paper the Q values for all actions are initialized to 0.

The agents update their expected gain in value for taking particular actions using equation 3, but how do they select an action to take? We follow the random meeting model of Watts (2001) for determining which two agents are going to consider their connection. Then, assuming that the two agents selected are i and j , if $ij \in G$ the agents will consider removing their connection. They will remove the edge if

$$\text{delete} = \arg \max_{a \in A} Q_i(a) \quad \text{or} \quad \text{delete} = \arg \max_{a \in A} Q_j(a). \quad (4)$$

That is, if either agent has as its highest value action to delete connections, the connection will be unilaterally deleted. The situation is similar for creating new connections, although the decision must be bilateral. That is, if $ij \notin G$, then the new connection will be created if

$$\text{add} = \arg \max_{a \in A} Q_i(a) = \arg \max_{a \in A} Q_j(a). \quad (5)$$

As is traditional in Q-learning, we allow the strategy to be ϵ -greedy (decisions are reversed with probability ϵ) with ϵ slowly decreasing with time.

The agents learn to switch between creating new connections, deleting connections, and leaving their current network structure unmodified. The strategy does not require any global computation to make decisions, although it does require the feedback signal for tracking changes that result from taking actions (or doing nothing). The key difference between the global computation method and our learning method is that we use past performance to determine future behavior rather than using a global calculation to determine what the future would be like. Furthermore, our strategy does not require that agents explicitly model other agents or the utility of any specific connections.

Results

Figure 2(a) shows the performance of our learning mechanism for various learning rates. As can be seen in the figure, the performance of our learning method is strongly dependent on the learning rate. If the learning rate is too slow, the agents tend to continue adding and deleting connections well beyond the time when those actions are beneficial. This leads to an oscillation in the value of the network, and to ultimate demise for the networked organization (e.g., the organization does not appear to recover from the gradual decrease in value for $\alpha = 0.05$). In our experiments, we found that higher learning rates performed much better, although agents using higher learning rates could be too quick to learn preventing the discovery of potentially high value structure. To alleviate this, we implemented an adaptive learning mechanism.

	average network value
global calculation	6955.82 \pm 8.512
adaptive local Q	7026.81 \pm 7.548

Table 1: The average performance of the adaptive learning rate network formation strategy compared with the average performance of the global computation strategy (25,000 iterations averaged over 10 simulations with 95% confidence intervals).

Adaptive Learning Rate. Following the method proposed for the “Win or Loss Fast” (WoLF) concept from Q-learning (Bowling & Veloso 2002), we incorporated an adaptive learning rate into our learning algorithm. As can be seen in Figure 2(a), the adaptive learning rate method outperforms all of the fixed rate learning schemes. The intuition behind the adaptive method is that an agent should learn “cautiously” when improving its position and an agent should learn rapidly when its decisions are resulting in decreases in value.

Using the adaptive learning rate, α_i is determined by

$$\alpha_i = \begin{cases} \alpha_{min} & \text{if } \Delta Y_i(G) > 0, \\ \alpha_{max} & \text{if } \Delta Y_i(G) < 0. \end{cases} \quad (6)$$

Using this adaptive scheme prevents an agent from getting “overexcited” about the gains received from taking any particular action, and allows an agent to learn rapidly when a certain action is detrimental. In the experiments presented in this paper, $\alpha_{min} = 0.05$ and $\alpha_{max} = 0.4$.

Looking at Figure 2(a), the adaptive learning rate strategy has performance similar to that of the global computation strategy. They both quickly form a highly connected network achieving a network value of approximately 7000. Both of the strategies plateau near a network value of 7000, as no single agent is willing to bear the greater burden of becoming the hub node in a star network structure. While their performance seems similar, the resulting network structures differ, and the difference in their average performance is statistically significant (although small in relative value). Table 1 shows the average value of the networks after 25,000 iterations over 10 simulations for the two strategies.

Figure 2(b) and 2(c) further suggest that the network structure found by the learning strategy differs from the structure that formed as a result of the global computation strategy. The network formed by the learning strategy is more dense, and contains a much greater amount of clustering (perhaps, as a direct result of the increased density). Another noteworthy result is that the cumulative number of rewirings flattens (bottom left of Figure 2(b)) for the learning strategy, where the global computation strategy continues to rewire. This is a result of our learning mechanism including the expected value of leaving the local connectivity structure unmodified. Eventually, the agents learn that the change in value for no modification (even if it is near zero) is greater than that of modifying its connections.

At this point, we conjecture that the increase in performance for our learning algorithm over that of the global computation strategy is a result of the inherent randomness




Figure 2: The result of distributed network formation when the agents employ the local learning strategy when $\delta = 0.9$ and $c = 0.8$ for a range of learning rates α : a) the value of the network over time, b) various structural properties of the network as it evolves, and c) the resulting network topology after 100,000 iterations. The adaptive strategy is the variable learning strategy discussed in the text.

in our method. The agents do not keep track of which connections lead to increased or decreased values, but rather, only the expected values of modifying connections. Therefore, when it comes time to add or delete a connection, the connection to modify is chosen as a result of the random meeting mechanism. We hypothesize that this allows the learning mechanism to do more exploration of the efficiency landscape.

Adding an Unselfish Agent

Any good decentralized network formation strategy should find the optimal network structure (i.e., the star) in the presence of a single unselfish agent. Therefore, we conducted an experiment where one of the agents was designated as completely unselfish. Let this unselfish agent be agent u . In order to force agent u to be unselfish, we set its Q values to be fixed at $Q_u(\text{add}) = 1.0$, $Q_u(\text{delete}) = 0.0$, and $Q_u(\text{nothing}) = 0.0$. Additionally, when u randomly meets another agent, it always convinces the other agent to establish a connection.

The results obtained from one of these experiments are shown in Figure 3. Our simple learning strategy is able to find a near optimal network structure in the presence of an unselfish agent. As before, the network quickly becomes highly connected and the network value levels off around 7000. After that, the agents learn that it is useful to prune away connections with any agents other than u . The small jumps (dips) in performance are a result of accidental deletion of connections between agents and the unselfish agent u .

Related Work

The formation of cooperation and interaction structures have been widely studied in the multi-agent systems literature. Coalition formation is the problem of identifying the optimal (or near-optimal) partition of the agents into groups so that the groups can work on specific tasks (Kraus, Shehory, & Taase 2003). Similar to coalition formation, congregation formation is the process of dividing a large collection

of agents into disjoint special interest groups for more efficient computation (e.g., agents with specific interest in a certain good should congregate in a specialized market for that good, reducing the computational cost of the auction for that good) (Brooks & Durfee 2003). Closely related to coalition formation, Dutta and Sen (2003) developed an agent-level learning mechanism that allowed the agents to learn which other agents were helpful, thus creating stable partnerships for task completion. The primary difference between the work on coalition and congregation formation and our work is the dependence of local and collective utility on the network in our work and our explicit concern for the structure of the agent social network.

There are several studies in the multi-agent systems literature that are more directly concerned with the formation of agent social networks. Specialized network adaptation mechanisms have been applied to information retrieval in peer-to-peer systems (Yolum & Singh 2003) and to dynamic team formation (Gaston & desJardins 2005). Additionally, several studies have considered models very similar to the SCM, including a unilateral connections model for analyzing the formation of large-scale communications networks (Fabrikant *et al.* 2003) and an extension of the SCM to a situation where the agents are spatially situated (Carayol & Roux 2004). Our work extends this previous work by developing and applying a simple (and potentially general) learning approach to dynamic, endogenous network formation.

Conclusion and Future Directions

We reviewed the theory and challenges of a simple network formation game, the Symmetric Connections Model, and discussed distributed network formation strategies. In addition, we developed a decentralized learning mechanism for network formation and demonstrated through simulation that distributed network formation is possible in the absence of a global computation.

There are many future directions for the continuation of this work. It would be interesting to study variations of the




Figure 3: The result of distributed network formation when the agents employ the local learning strategy when $\delta = 0.9$ and $c = 0.8$ with an adaptive learning rate α in the presence of exactly one unselfish agent: a) the value of the network over time, b) various structural properties of the network as it evolves, and c) the resulting network topology after 50,000 iterations.

dynamic network formation process of Watts (2001) by having the agents make decisions based on local calculations (i.e., the change in value in the local neighborhood) and by going beyond the myopic decision criteria. More generally, we would like to move beyond the random meeting model for determining the pair of agents that will consider their connection, and move toward a simultaneous, asynchronous network formation process (i.e., where all of the agents can make decisions simultaneously). Finally, we hope to apply the theory and findings for the simple SCM to more complicated networked multi-agent environments, such as distributed production and exchange, team formation, supply chain formation, and distributed information retrieval.

References

- Bowling, M., and Veloso, M. 2002. Multiagent learning using a variable learning rate. *Artificial Intelligence* 136(2):215–250.
- Brooks, C. H., and Durfee, E. H. 2003. Congregation formation in multiagent systems. *Journal of Autonomous Agents and Multiagent Systems* 7(1-2):145–170.
- Carayol, N., and Roux, P. 2004. Behavioral foundations and equilibrium notions for social network formation processes. *Advances in Complex Systems* 7(1):77–92.
- Claus, C., and Boutilier, C. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 746–752.
- Culler, D.; Estrin, D.; and Srivastava, M. 2004. Overview of sensor networks. *IEEE Computer* 37(8):41–19.
- Dutta, P. S., and Sen, S. 2003. Forming stable partnerships. *Cognitive Systems Research* 4(3):211–221.
- Fabrikant, A.; Luthra, A.; Maneva, E.; Papadimitriou, C. H.; and Shenker, S. 2003. On a network creation game. In *Proceedings of the 22nd Annual ACM Symposium on the Principles of Distributed Computing*, 347–351. ACM Press.
- Gaston, M., and desJardins, M. 2005. Agent-organized networks for dynamic team formation. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*.
- Jackson, M. O., and Watts, A. 2002. The evolution of social and economic networks. *Journal of Economic Theory* 106:265–295.
- Jackson, M. O., and Wolinsky, A. 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71:44–74.
- Jackson, M. O. 2003. A survey of models of network formation: Stability and efficiency. In Demange, G., and Wooders, M., eds., *Group Formation in Economics: Networks, Clubs, and Coalitions*. Cambridge University Press.
- Kraus, S.; Shehory, O.; and Taase, G. 2003. Coalition formation with uncertain heterogeneous information. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '03)*.
- 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*.
- Thadakamalla, H. P.; Raghaven, U. N.; Kumera, S.; and Albert, R. 2004. Survivability of multiagent-based supply networks: A topological perspective. *IEEE Intelligent Systems* 19(5):24–31.
- Walsh, W. E., and Wellman, M. P. 2003. Decentralized supply chain formation: A market protocol and competitive equilibrium analysis. *Journal of Artificial Intelligence Research* 19:513–567.
- Watts, A. 2001. A dynamic model of network formation. *Games and Economic Behavior* 34(331-341).
- Yolum, P., and Singh, M. P. 2003. Dynamic communities in referral networks. *Web Intelligence and Agent Systems (WIAS)* 1(2):105–116.
- 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.