Agent-Organized Networks for Multi-Agent Production and Exchange

Matthew E. Gaston and **Marie desJardins**
Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
{mgasto1, mariedj}@cs.umbc.edu

Abstract

As multi-agent systems grow in size and complexity, social networks that govern the interactions among the agents will directly impact system behavior at the individual and collective levels. Examples of such large-scale, networked multi-agent systems include peer-to-peer networks, distributed information retrieval, and agent-based supply chains. One way of dealing with the uncertain and dynamic nature of such environments is to endow agents with the ability to modify the agent social network by autonomously adapting their local connectivity structure. In this paper, we present a framework for agent-organized networks (AONs) in the context of multi-agent production and exchange, and experimentally evaluate the feasibility and efficiency of specific AON strategies. We find that decentralized network adaptation can significantly improve organizational performance. Additionally, we analyze several properties of the resulting network structures and consider their relationship to the observed increase in organizational performance.

Introduction and Related Work

The success of both real and artificial organizations is dependent upon a structure that facilitates effective and efficient behavior at the individual and organizational levels. In many multi-agent system applications, groups of agents must coordinate to solve problems, efficiently distribute goods or services, form teams to accomplish tasks, and collect and share information. In these domains, the organizational structure, or the agent social network, has a direct impact on the performance of the agent society. Our goal is to enable agents to autonomously adapt their local network connectivity, providing organizations with a level of social intelligence and an organizational learning capability.

In this paper, we develop an approach for implementing agent-organized networks (AONs) in the context of a multi-agent production and exchange economy (Wilhite 2003). Potential application domains for this work include the management and formation of supply chain networks (Geunes & Pardalos 2003; Fox, Barbuceanu, & Teigen 2000), peer-to-peer (P2P) networks (Ramanathan, Kalogeraki, & Pruyne 2002), and distributed information retrieval (Yu & Singh 2003). Researchers have recently emphasized the necessity of decentralized, bottom-up supply chain formation:

> "To achieve the oft-expressed visions of dynamically forming and dissolving business interactions requires automated support for *supply chain formation*, the process of bottom-up assembly of complex production and exchange relationships." (Walsh & Wellman 2000)

Related to our work in economically motivated AONs, network formation has been studied in economic game theory (Jackson 2003), including applying reinforcement learning to refine partner selection in repeated games (Skyrms & Pemantle 2000). Small-scale trade networks and partner selection have been studied in agent-based computational economics (Tesfatsion 1997), and reinforcement learning has been proposed as a way to learn effective agent interactions based on reputation in multi-agent market environments (Tran & Cohen 2003). Dutta and Sen examined learning cooperative relationships for efficient task completion using a "simple reinforcement scheme" (2003).

Our work extends and differs from the previous research in two important ways. First, we take a bottom-up approach to network formation, maintaining the resource, cognitive, and communication limitations implied by an initial agent social network. Second, we explicitly analyze the structure of the agent social network resulting from various decentralized adaptation strategies. Two assumptions guiding our approach are the lack of a central broker or global search capability and the assumption that the agents may not be under the control of a single authority.

The goals of this paper are to provide a framework for bottom-up AONs and to demonstrate the feasibility of applying AONs to improve the performance of an organization of economically motivated agents. We are particularly interested in developing agent-initiated network adaptation strategies that are realistic, feasible, and efficient. We first describe a general multi-agent production and exchange model and the AON framework. We then present experimental results demonstrating the feasibility of AONs for multi-agent production and exchange.

Multi-Agent Production and Exchange

In order to study mechanisms for AONs in multi-agent economies, we selected a simple, generic, yet realistic model of a production and exchange economy. In this section, we describe the model, and discuss the effects of agent social structures on the organization's ability to distribute goods.

A Model of Production and Exchange

The basis for the model was first presented by Wilhite (2001; 2003). Each agent is given an initial endowment of two distinct goods, and has a fixed production capacity. At each time step, each agent is allowed to choose whether to trade or to produce. The goal of the individual agents is to maximize their utility. The model assumes that agents are purely selfish (i.e., they select the action that maximizes their utility) and completely truthful (i.e., they always provide perfect information during negotiation and trade).

Let there be n agents in the economy and two goods, g_1 and g_2, where g_2 is infinitely divisible and g_1 must be traded in whole units.[1] The utility of agent i is

$$U^i = g_1^i g_2^i. \quad (1)$$

It follows that if an agent possesses a total of $G = g_1 + g_2$ goods, then the optimal allocation is $g_1 = g_2$.

In the original model, the agents were given the ability to produce a set amount of both goods. In order to promote trading among the agents, we restrict the agents to being able to produce one good. This restriction requires that agents trade to maximize utility. Assuming that agent i is a producer of g_1, $\Delta g_1^i \in [1, q]$ and $\Delta g_2^i = 0$ (and likewise for producers of g_2) where q is a model parameter. This allows for a society of heterogeneous agents, in which some are effective producers and others are poor producers. The latter must rely more heavily on trade to increase their utility.

At each iteration, the agents are selected in random order and are allowed to negotiate (i.e., determine the price of trading) with m other agents. The selected agent then chooses the action—trade with one of the m agents or produce—that maximizes its utility. In negotiation, each agent truthfully reveals their *marginal rate of substitution*,

$$mrs_i = \frac{\delta U^i / \delta g_1^i}{\delta U^i / \delta g_2^i} = \frac{g_2^i}{g_1^i}. \quad (2)$$

When the two negotiating agents' marginal rates of substitution differ, there is an opportunity for mutually beneficial trade between the two agents (Wilhite 2003). Assuming that agent i is negotiating with agent j, the next step in the negotiation process is to calculate the exchange price:

$$p_{i,j} = \frac{g_2^i + g_2^j}{g_1^i + g_1^j}. \quad (3)$$

The agents must also take into account the trading tax τ, which is given as a model parameter. Assuming that agent i is trading one unit of g_1 for $p_{i,j}$ units of g_2 with agent j, the tax is applied to the transaction such that

$$g_1^i = g_1^i - (1.0 + \tau) \quad \text{and} \quad g_2^j = g_2^j - (1.0 + \tau) p_{i,j}. \quad (4)$$

[1]Forcing g_1 to be traded in whole units is to simplify the price formation and trading process. It is claimed that this adds realism to the model (Wilhite 2003).

The agents repeatedly trade in this manner until the exchange no longer increases the utility of either agent. During negotiation, the agents do not actually exchange goods; rather, the active agent computes the change in utility $\Delta U^i(j)$ that would result from trading with each of the m agents with whom it has negotiated. The agent also calculates its change in utility after production:

$$\Delta U^i(g_1) = (g_1^i + \Delta g_1^i) g_2^i, \quad (5)$$

assuming that agent i is a producer of g_1. Finally, the agent selects the action that results in the largest ΔU^i and then all of the goods are either exchanged or produced accordingly.

This model was originally studied in the context of global price conversion and the roles that individual agents adopted: heavy traders, heavy producers, and specialized producers. The nature of the interactions in the initial study was based on each agent selecting m other agents at random to interact with at each time step (Wilhite 2003). Here, we refer to this interaction paradigm as *random mixture*, and use it as a baseline for performance in our experiments.

An alternative to random mixture is to embed the agents in a fixed network topology and only allow agents that are directly connected to negotiate and trade. It has been shown that network structure has a direct impact on the rate at which the economy converges on a global price (Wilhite 2001) when the agents only trade (i.e., have no production capacity). In the next section, we demonstrate this dependence using a larger set of network structures and include the production capabilities of the agents.

Effects of Trading Structure

In order to evaluate the efficiency of the trading network, we introduce a global measure for the *skew* of goods,

$$S = \frac{\sum_i (g_1^i - g_2^i)^2}{\sum_i (g_1^i + g_2^i)^2}. \quad (6)$$

Intuitively, S measures the total imbalance of the two goods for all of the agents, relative to the total amount of goods in the system. Since the agents in our model can only produce one of the two goods, trade is the only way for an agent to balance its distribution of goods (i.e., move toward $g_1^i = g_2^i$).

To motivate the need for AONs, we measured skew after 500 iterations of the model for each of five static network structures: **one-dimensional lattices** (agents organized in a ring, connected only to their nearby neighbors), **two-dimensional lattices** (agents organized in a two-dimensional toroidal grid), **one-dimensional small-worlds** (one-dimensional lattice with a small percentage of connections randomly rewired to provide "short-cut links" (Watts & Strogatz 1998)), **random graphs** (connections exists between agents i and j with probability p (Erdos & Renyi 1959)), and **star topologies** (hub agents are connected to all other agents in the system, approximating a scale-free network (Albert & Barabási 2002)).

Figure 1 shows the average skew for 25 trials for the five network structures as a function of τ. (The error bars are 95% confidence intervals.) All of the network structure were parameterized to have the same number of nodes (i.e.,

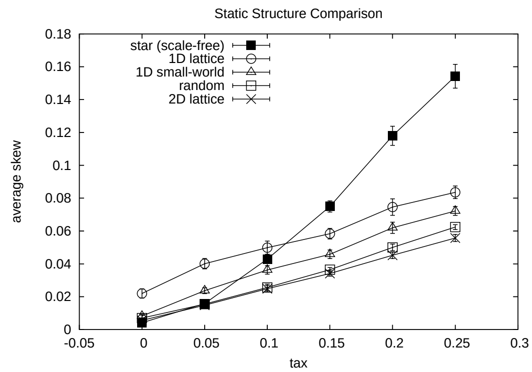

Figure 1: Average skew as a function of the trading tax for various static network structures.

agents) and connections as the two-dimensional lattice (400 agents and 800 undirected connections).

As expected, there is a statistically significant difference in the resulting skew for the different network structures in line with previous results (Wilhite 2001). Interestingly, the star topology yields the lowest skew when there is no tax on trading since the central agents, or hubs, serve as "clearinghouses" for the goods produced by the agents on the periphery of the network. As the tax is increased, these central agents can no longer "afford" to serve as clearinghouses.

The dramatic effect that the network structure has on the performance of the production and exchange economy motivates our research on decentralized mechanisms that allow individual agents to adapt their local network connectivity.

AONs in the Multi-Agent Economy

The term *agent-organized network* (AON) refers to an adaptive agent interaction topology where individual agents change their local connectivity in real time. AONs are well suited for dynamic agent environments in which multiple authorities are responsible for the development and deployment of the agents. They are also useful when agents only have access to local, and potentially uncertain, information.

Compared to centralized network adaptation methods, the design of AON mechanisms presents a set of unique challenges. First, the agents must base their decisions on local, incomplete information. This can lead to poor adaptation decisions that will need to be corrected in the future. AONs also involve decentralized and simultaneous adaptation, which could negate the benefit of adaptation. That is, if two "nearby" agents in an organization decide to adapt their local connectivity simultaneously, they could both end up in worse positions in the network, even though each adaptation alone may have led to an improvement.

We are only concerned with AON strategies that preserve the resource, cognitive, and communication constraints of the starting network topologies. Therefore, we do not allow for a change in the number of connections in the network: agents can only *rewire* (i.e., they must drop a connection to one neighbor and simultaneously add a connection to a different agent). This allows for unbiased comparison between static network structures and adapted network structures resulting from AON strategies.

Our current approach is inherently bottom-up: we start with an existing network structure and attempt to bootstrap the structure into a more efficient organization. This is in direct contrast with previous top-down approaches, which assume that all agents can initially interact with all other agents; over time, the agents learn (generally through some form of reinforcement) to prefer interacting with certain other agents (Skyrms & Pemantle 2000; Wilhite 2003; Tran & Cohen 2003). While the top-down approach is reasonable, it does assume a global search mechanism which is not realistic for many applications.

An AON Framework

Our AON framework consists of three components: deciding when to adapt, selecting which connections to rewire, and selecting agents for new connections.

Deciding When to Adapt The intuition here is that an agent will not change its local network connectivity as long as the agent maintains a certain level of positive change in utility. Let V_t^i be the *expected change in utility* for agent i at time t. We update the value of an agent's expected change in utility using an exponentially weighted moving average:

$$V_{t+1}^i = V_t^i + \alpha(\Delta U_t^i - V_t^i), \tag{7}$$

where ΔU_t^i is the change in utility at the current time step and α is the learning rate. In the context of the production and exchange economy, the expected change in utility is the expected level of growth and ΔU_t^i is the resulting change in utility following all trading (initiated by any agent) and production. Modeling the expected increase in utility in this way, an agent will chooses to adapt when $V_t^i < \Theta$.

Deciding Which Connections to Rewire This decision is also based on an exponentially weighted moving average. Each agent maintains a *connection value* for each of its local connections, updated using the equation:

$$W_{t+1}^{i,j} = W_t^{i,j} + \beta(T_t^{i,j} - W_t^{i,j}), \tag{8}$$

where $T_t^{i,j}$ is the actual value of the connection at time t and β is the connection learning rate. In the production and exchange economy, $T_t^{i,j}$ is the number of trades made using that connection at time t. The maximum value for $T_t^{i,j}$ is 2, since both i and j could possibly initiate trade along this connection at a given time step as the connections are undirected. When an agent decides to adapt its local network structure, it rewires each connection if $W_t^{i,j} < \Phi$.

Deciding Where to Make A New Connection This is perhaps the most challenging part of the framework. There are many candidate approaches for determining the destination agent of a new connection. One approach, *random selection*, is to select a new agent to connect with at random. Similar to random mixture, random selection implies the existence of a global search capability.

Another approach for establishing new connections is *referrals*. Referrals are more realistic than random searches in a decentralized environment, but referrals come with their own set of challenges. One such challenge is the possibility of disconnecting the network (i.e., breaking the network into disjoint components). The potential for referrals to disconnect the network is problematic as referrals cannot reconnect a disconnected network. To avoid this problem, we develop a special class of referral-based adaptation strategies.

AON Strategies

We present three AON strategies based on the notion of *push referrals*. For the remainder of the paper, we will use the notation $N_j(i)$ to represent the set of agents that are neighbors of agent j but not neighbors of agent i.

Definition 1 *Assuming that agent i is adapting its connection to agent j, a* **push referral** *is a local rewiring by i from j to an agent in $N_j(i)$.*

Although restrictive, push referrals have a highly intuitive interpretation, especially in cooperative multi-agent systems. Imagine going into a store to purchase a certain product. If the store is out of that product, it is reasonable to ask the salesperson for a recommendation of a nearby location that may have the product in stock. Besides its intuitive appeal, given an initially connected network of agents (i.e., one component), an AON that is a sequence of push referrals guarantees that the network will remain connected.

We propose three strategies for push referral AONs. The methods described in the previous section are used to determine when to adapt and which connections to rewire. The following three push referral heuristics are used to select the agent for the new connection.

- **Random Referral**: agent j randomly selects an agent $r \in N_j(i)$ to refer. The probability that i establishes a connection with agent r is $P(i \rightsquigarrow r) = 1/|N_j(i)|$.
- **Degree Referral**: agent j selects an agent $r \in N_j(i)$ to refer with probability proportional to the degree k_r of agent r. The probability that an agent is selected is

$$P(i \rightsquigarrow r) = \frac{k_r}{\sum_{l \in N_j(i)} k_l}. \quad (9)$$

  This referral strategy is based on our analysis of static networks and the finding that the star topology (under low trade tax τ) is an efficient structure. There is a large amount of research on the ubiquity of scale-free networks, for which the canonical formation model is preferential attachment by degree (Albert & Barabási 2002).
- **Production Referral**: agent j selects an agent $r \in N_j(i)$ to refer with probability proportional to agent r's ability to produce the complimentary good of agent i. The probability of a new connection between i and r is

$$P(i \rightsquigarrow r) = \frac{e^{\Delta g_2^r}}{\sum_{l \in N_j(i)} e^{\Delta g_2^l}}, \quad (10)$$

  assuming that agent i is a producer of g_1. This referral strategy is designed specifically for the structure of the model. While this strategy favors agents that produce a large amount of the complimentary good, the probability distribution allows for connections with agents that are not producers of the complimentary good (as $e^0 = 1$).

We next present experimental evaluations of these strategies and analyze the resulting network structures.

Experiments and Results

In this section, we provide a representative sample of our results. For the experiments, the model was parameterized following Wilhite (2003; 2001): $n = 400$, $q = 30$, $\tau = 0.05$, and the initial endowments were chosen uniformly at random from the interval $[1, 60]$. The AON learning and threshold parameters were set at $\alpha = \beta = 0.1$ and $\theta = \Phi = 0.1$. Finally, the agents' expected utilities, V^i, and the values of the connections, $W^{i,j}$, were all initialized to 1.0. Our findings for a wide range of parameter settings were qualitatively similar to the results presented in this section.[2] The parameter that affected performance of the AON strategies the most was the initial network topology (see Table 1).

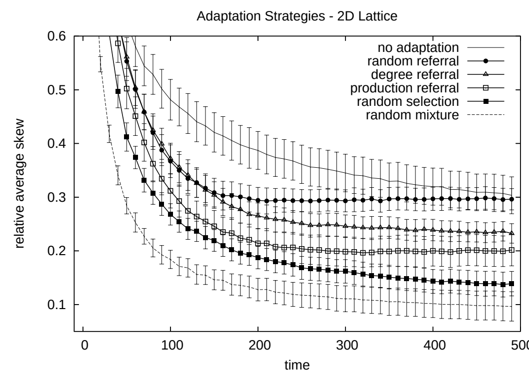

Figure 2: Average skew relative to the initial skew over time for the various AONs, random mixture, and no adaptation. The initial networks are two-dimensional lattices.

Figure 2 shows the relative average skew over time. We chose to present the results starting from two-dimensional lattices since they performed the best in our study of static networks. The three benchmarks are the lattice with *no adaptation*, *random mixture* with $m = 4$ (as the initial degree of every agent in the lattice is 4), and the *random selection* AON strategy. We include the random selection strategy as it provides an estimate of the theoretical potential of AONs. The data are the average of 25 simulations and the error bars are 95% confidence intervals.

There is a statistically significant decrease in skew over time for both the degree and production referrals compared with no adaptation and the random referral. This is expected

[2] We experimented with a range of learning parameters (i.e., $\alpha, \beta, \theta, \Phi \in [0.01, 0.2]$) and found only that the rate of learning was affected. Similarly, only the scale of the skew values and not the relative performance of the various strategies was affected for a range of values of q, the initial endowments, and τ. The initial values for V^i and $W^{i,j}$ had little effect on the performance as long the initial values had $V^i > \theta$ and $W^{i,j} > \Phi$.

since the degree and production referral strategies are exploiting the structure of the environment. Note that while their confidence intervals overlap, as time progresses, the production referral shows a slight advantage over that of the degree referral. Finally, the performance of random selection suggests that there remains room for improvement in the design of AON strategies.

	2D lattices	1D lattices	random graphs
rs	54.2 (7.53)	82.3 (2.84)	47.7 (10.9)
pr	33.6 (9.75)	58.1 (5.43)	1.12 (10.9)
dr	23.4 (5.99)	57.9 (4.66)	-13.2 (7.50)
rr	2.41 (6.53)	49.4 (4.98)	-41.7 (7.61)

Table 1: Percentage decrease in skew compared with no adaptation starting with two-dimensional lattices ($S = 0.015$), one-dimensional lattices ($S = 0.040$), and random graphs ($S = 0.014$). (*rs = random selection, pr = production referral, dr = degree referral, rr = random referral*.)

As mentioned above, the starting network has a significant effect on the performance of the adaptation schemes. Table 1 shows significant decrease in skew as a result of adaptation starting from one dimensional lattices. Alternatively, the utility of the referral-based adaptation strategies is significantly diminished when the initial network topologies are random graphs. We hypothesize that this is related to the lack of structure and the short average path length of random graphs. At the same time, random selection provides a significant increase over the static random graph demonstrating the utility of adaptation and reiterating the need for well designed AON strategies.

Comparing the push referral strategies to the non-referral strategies is somewhat misleading, since the current model does not factor the cost of searching and rewiring. It is important to factor the cost of searching, since the non-referral strategies require $O(n)$ searches for any adaptation, where the push referral adaptations require $O(1)$ searches (on average). For random mixture, every agent, at every time step, randomly selects m agents from the pool of $n-1$ agents; each agent effectively performs $n-1$ searches and m "rewirings" at each time step. In our experiments, $m = 4$ and $n = 400$, yielding $n(n-1) = 159,600$ searches and $nm = 1,600$ rewirings **during each time step**. In comparison, each of the AON strategies (including random selection) experimentally averaged approximately 3,600 rewirings over the entire 500 time steps. The push referral strategies averaged three searches for each rewiring (since the average number of connections of an agent is always four), resulting in only 10,800 total searches. Therefore, the push referral based AON strategies yield a 99.99% decrease in the number of searches and a 99.55% decrease in the number of rewirings compared to random mixture. Although random selection averaged the same number of rewirings as the push referral strategies, it requires significantly more searches. An agent using random selection must effectively search the entire organization for each new connection, with the exception of the agents with whom the adapting agent is currently connected. This results in $n - m - 1 = 395$ searches for each rewiring. Therefore, the push referral strategies decreases the number of necessary searches by 99.2% compared to random selection. As a result, push referral AON strategies will provide a more scalable solution for network adaptation over global search strategies when searching is costly.

Discussion

To better understand the behavior of the AON strategies, we collected structural statistics of the networks before and after AON adaptation. The mean path length D is the average shortest path between all pairs of agents in the network (Albert & Barabási 2002; Newman 2003). The clustering coefficient C is the ratio of transitivity (i.e., triangles) in the network. We use the localized calculation, $C = \frac{1}{n}\sum_i \frac{2|E_k^i|}{k_i(k_i-1)}$, where E_k^i is the set of connections among the k_i neighbors of agent i, and k_i is the degree of agent i (Newman 2003). The production–degree correlation P_k is the standard correlation between the degree of the agents in the network and the amount they produce (of either good). Finally, we measured the production correlation $\rho(\Delta g_1, \Delta g_2)$ of adjacent agents in the network using an adaptation of Newman's assortativity measure, which is a version of Pearson's correlation coefficient (Newman 2003). Note that $\rho \in [-1, 1]$, where a strong positive value means large producers of one good are more regularly adjacent to large producers of the opposite good, and a strong negative value suggests that large producers of one good are more regularly adjacent to weak producers of the opposite good.

The table in Figure 3(a) shows the average values of these statistics for 25 simulations starting from two-dimensional lattices. On the left of Figure 3(b) is a rendering of the initial two-dimensional toroidal lattice. On the right is an example of a network resulting from the production referral strategy. The table shows that the initial network structure has an average path length of 10 and, as expected, no clustering, production-degree correlation, or production correlation. All of the adaptation strategies significantly reduce the average path lengths between the agents in the networks with the referral strategies yielding shorter path lengths than random selection. Similarly, the referral strategies significantly increase clustering, where random selection produces no more clustering than would be expected in a random graph of the same size (i.e., 0.01). An increase in P_k suggests that the agents that are large producers tend to become hub agents, or more centralized. The production referral strategy's ability to significantly increase both P_k and ρ may be indicative of its slight advantage over the other referral strategies, although it was unable to outperform random selection. Finally, random referral does not yield as effective of a trading structure, which may be a result of its overemphasis on increasing clustering and inability to increase P_k.

Although the referral strategies underperform random selection, they do create structures for efficient trading. Importantly, they do this with far less searching and no global search capability. The performance results, coupled with the structural statistics of the resulting networks, suggest that push referral based AONs are an effective decentralized

Resulting Network Structural Statistics				
	D	C	P_k	ρ
na	10 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)
rs	6.0 (0.86)	0.01 (0.00)	0.18 (0.03)	0.06 (0.02)
pr	4.4 (0.35)	0.06 (0.01)	0.34 (0.02)	0.28 (0.02)
dr	4.0 (0.30)	0.08 (0.02)	0.16 (0.02)	0.03 (0.02)
rr	4.4 (0.29)	0.17 (0.02)	0.04 (0.02)	0.02 (0.03)

(a)

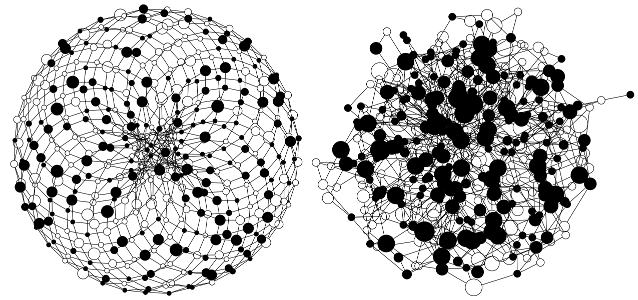

(b)

Figure 3: Network structures: a) table of structural statistics before and after adaptation starting from two-dimensional toroidal lattices (*na = no adaptation*, *rs = random selection*, *pr = production referral*, *dr = degree referral*, and *rr = random referral*), and b) the starting two-dimensional lattice (left) and a network resulting from production referral (right) where the size of a node is proportional to its production capacity and its fill denotes the good it produces (producers of g_2 are filled).

method of discovering efficient trading structures.

Conclusions and Future Work

We have presented a framework for AONs and demonstrated their feasibility for improving the efficiency of a general, networked multi-agent economy. We have shown that push referral AON strategies, which require no global search capability, can discover efficient trading structures.

Our ongoing research is focused on developing theoretical aspects of AONs and on developing additional decentralized rewiring strategies. In the near future, we hope to develop hybrid strategies, that may allow the network to become disconnected and subsequently reconnected. We are exploring variable learning rates and methods for incorporating the states of neighboring agents into an agent's decision making to provide a more direct multi-agent learning approach. Finally, we are planning to employ AONs in a wide variety of multi-agent environments.

Acknowledgments

We thank Allen Wilhite, Qianjun Xu, Blazej Bulka, Eric Eaton and the anonymous reviewers for their valuable comments and suggestions.

References

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

Dutta, P. S., and Sen, S. 2003. Forming stable partnerships. *Cognitive Systems Research* 4(3):211–221.

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

Fox, M. S.; Barbuceanu, M.; and Teigen, R. 2000. Agent-oriented supply-chain management. *International Journal of Flexible Manufacturing Systems* 12(2-3):165–188.

Geunes, J., and Pardalos, P. M. 2003. Network optimization in supply chain management and financial engineering: An annotated bibliography. *Networks* 42(2):66–84.

Jackson, M. O. 2003. A survey of network 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.

Newman, M. 2003. The structure and function of complex 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)*.

Skyrms, B., and Pemantle, R. 2000. A dynamic model of social network formation. *Proceedings of the National Academy of Sciences USA* 97(16):9340–9346.

Tesfatsion, L. 1997. A trade network game with endogenous partner selection. In Amman, H. M.; Rustem, B.; and Whinston, A. B., eds., *Computational Approaches to Economic Problems*. Kluwer Academic Publishers.

Tran, T., and Cohen, R. 2003. Learning algorithms for software agents in uncertain and untrusted market environments. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*.

Walsh, W. E., and Wellman, M. 2000. Modeling supply chain formation in multiagent systems. In *Lecture Notes in Artificial Intelligence: Agent-Mediated Electronic Commerce II*, volume 1788. Springer-Verlage.

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

Wilhite, A. 2001. Bilateral trade and 'small-world' networks. *Computational Economics* 18(1):49–64.

Wilhite, A. 2003. Self-organizing production and exchange. *Computational Economics* 21(1-2):107–123.

Yu, B., and Singh, M. P. 2003. Searching social networks. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. ACM Press.