

Agent-Organized Networks for Dynamic Team Formation

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

ABSTRACT

Many multi-agent systems consist of a complex network of autonomous yet interdependent agents. Examples of such networked multi-agent systems include supply chains and sensor networks. In these systems, agents have a select set of other agents with whom they interact based on environmental knowledge, cognitive capabilities, resource limitations, and communications constraints. Previous findings have demonstrated that the structure of the artificial social network governing the agent interactions is strongly correlated with organizational performance. As multi-agent systems are typically embedded in dynamic environments, we wish to develop distributed, on-line network adaptation mechanisms for discovering effective network structures. Therefore, within the context of dynamic team formation, we propose several strategies for agent-organized networks (AONs) and evaluate their effectiveness for increasing organizational performance.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multi-agent Systems;
I.2.11 [Distributed Artificial Intelligence]: Coherence and Coordination

General Terms

Algorithms

Keywords

Team Formation, Multi-agent Systems, Organizational Learning

1. INTRODUCTION

In both real and artificial societies, successful organizations are highly dependent upon a structure that fosters effective and efficient behavior at both the individual and the organizational levels. In many applications of multi-agent systems, groups of agents must coordinate effectively in order to solve problems, allocate tasks across a distributed organization, collectively distribute knowledge and information, and achieve collective goals. The organizational

structure of a multi-agent system dictates the interactions among the agents, and can play a significant role in the overall performance of the agent society.

Recently, several studies have emphasized the importance of the social structure of agent systems and the impact that structure has on organizational performance. The rate of the spread of social conventions in a society of agents is affected by the complex network structure among the agents [10]. Line-of-sight considerations, restricting communication and interaction, affect the performance of mobile robot teams [5]. 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 [7]. In the context of multi-agent team formation, networks that exhibit short average path length [3, 23, 30] and a hub structure [3, 23] allow for greater diversity in teams of agents and greater efficiency in forming teams [13, 14]. Finally, recent research on multi-agent games in complex network structures demonstrated that variations in agent social networks lead to variance in the stability of agent strategies [2, 15, 18, 27].

Given that network structures have a dramatic effect on distributed agent systems, we are interested in the practicality of developing decentralized mechanisms that allow individual agents to manage their local network connectivity. Understanding how the interaction topology of an agent system enables or restricts organizational efficiency motivates the need for agent-driven network adaption and guides the development of agent-organized networks.

There are some important real-world applications of AONs, including sensor networks and supply chain management. In these applications, agents must deal with dynamic and uncertain environments. In addition, it is possible that the design and control of the agents do not fall under a single authority.

Sensor networks are rapidly becoming an important area of research for the multi-agent systems community. Applications of sensor networks include environmental monitoring, structural modeling, disaster management, health care, and manufacturing [8]. In particular, wireless sensor networks present several unique challenges, including network connectivity among sensors (i.e., agents) situated in some physical space. Although the physical space largely determines the connectivity of a sensor network, there are aspects of connectivity over which the agents have control, or to which they must adapt. Many different types of agents can be included in sensor networks, including sector managers, data collectors, data routers, and end point sensors. In these situations, the role that an agent takes on and the interconnectivity of the agents is important for the overall efficiency of the network. Strategies and policies for intelligently adjusting network connectivity among heterogeneous agents in a sensor network may lead to improved performance and fault tolerance [8, 16].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

Supply chain management is another important multi-agent application domain that can benefit from the study of AONs [6, 12]. The topology of a multi-agent based supply network has a direct influence on the efficiency and robustness of the system [29]. Furthermore, the components of a supply network typically do not fall under the control of a single authority. Endowing the agents with the capability to systematically manage their connectivity in a supply network will allow for robust performance in dynamic supply and demand environments.

The goal of this paper is to assess the feasibility of designing realistic and efficient strategies for distributed network adaptation for improving organizational performance. In the next section, we introduce and discuss the concept of agent-organized networks.

2. AGENT-ORGANIZED NETWORKS

It is common practice to design or specify the organizational structure of a multi-agent system. An alternative is to design local network adaptation strategies for individual agents that give rise to dynamic, adaptable organizational structures. We introduce the term *agent-organized network* (AON) to refer to such structures:

DEFINITION 2.1. *An agent-organized network (AON) is an organizational network structure, or agent-to-agent interaction topology, that is the result of local rewiring decisions made by the individual agents in a networked multi-agent system.*

Here, *networked multi-agent system* refers to a multi-agent system where there is an explicit set of interactions, or interdependencies, among the agents in the system. AONs are particularly well suited for agent environments where multiple parties are responsible for the development and deployment of the agents. AONs are also applicable in dynamic environments where the agents have only local and possibly uncertain information.

The design of AONs presents a set of unique challenges. First, agents must make network adaptation decisions based on local information about the agent system. This local information could give the agent a partially or completely incorrect picture of the actual environment. In this setting, agents have a limited horizon and may make adaptation decisions that are detrimental to the agent or the organization. In addition, because AONs involve decentralized and asynchronous adaptation, simultaneous decisions could negate the benefit of adaptation. That is, if two “near-by” agents in the organization decide to adapt their local connectivity at the same time, they could both end up in worse positions in the organization than prior to adaptation.

We are particularly focused on AON strategies for cooperative multi-agent systems. In this scenario, the goal of the individual agents is to increase the collective performance of the agent organization. The AON strategies of concern in this paper are limited to *rewiring* of an initial, arbitrary network topology. That is, the number of connections among the agents in an organization is fixed, and adaptation can only occur through the one-for-one removal and addition of a connection. While AONs could employ explicit creation and removal of connections, we are focused solely on rewiring strategies in order to preserve resource, cognitive, and communications constraints based on the initial network topology.

Important design considerations for AOC strategies in cooperative environments include:

- **Local perception of global performance:** How does an agent estimate the collective performance of the organization? These estimates will be uncertain, since they are based on partial observations of the organization. Communication capabilities determine whether the agent forms an estimate

independently or if the agent can incorporate estimates of its neighbors.

- **Adaptation triggers:** When should an agent decide to adapt its local connectivity structure? There are many possibilities for triggering network adaptation. An agent could decide to adapt based on performance estimates; structural requirements, such as neighborhood diversity, could drive adaptation; or simple timing mechanisms could be utilized.
- **Rewiring:** How does an agent decide on a connection to remove and how does an agent select a new agent with whom to establish a connection? There is a seemingly infinite number of possibilities for specifying rewiring. Referrals based on various criteria and decisions based on historical behavior are two of the many possibilities.

There are also several important issues to consider in evaluating the effectiveness of AON strategies:

- **Learning rate:** How quickly does a particular adaptation strategy increase (or decrease) performance?
- **Stability:** Does the adaptation strategy lead to a stable (or possibly periodic) network topology? Is stability beneficial or detrimental to organizational performance?
- **Global structure:** What are the properties of the network structures that result from AON strategies? How do these structural properties evolve over time?

In order to begin the exploration of AONs, we have selected a simple, but intuitive, model of a cooperative multi-agent system. In the next section, we briefly describe an environment for networked multi-agent team formation. We then introduce two AON strategies for this domain, conduct experiments to evaluate their utility for improving collective team formation performance, and discuss their behavior.

3. DYNAMIC MULTI-AGENT TEAM FORMATION

To explore agent-organized networks in a dynamic environment, we have selected a simple multi-agent system model motivated by previous work on agent team formation [1, 14, 19, 21]. The model provides a dynamic team formation environment where agent teams form spontaneously in a completely decentralized manner and the agents’ decision making is based solely on local information. In addition, the model allows for potentially very large agent organizations where the agents are embedded in an agent social network. The model is only concerned with the dynamic formation of teams and is not concerned with teamwork mechanisms or protocols, for which there is a large body of previous work [17, 20, 22, 24, 28].

In the model, tasks are generated periodically and globally advertised to the organization. Agents attempt to form teams to accomplish these tasks. The agents in the organization are embedded in a social network that restricts the set of possible agent teams: specifically, for an agent to be on a team, the agent must have a social connection (i.e., an edge in the social network) with at least one other agent on the team. Since we are only concerned with the formation process, tasks are generic in that they only require that a team of agents with the necessary skills form to accomplish the specific task.

In this model of team formation, the organization consists of N agents, $A = \{a_1, a_2, \dots, a_N\}$, where each agent can be considered as a unique node in the social network. The network is

modeled as an adjacency matrix E , where an element of the adjacency matrix $e_{ij} = 1$ if there is an edge between agent a_i and a_j and $e_{ij} = 0$ otherwise. The social relationships among the agents are undirected, so $e_{ij} = e_{ji}$, and for all agents, $e_{ii} = 0$. The *degree* (number of connections) of agent a_i is defined as $k_i = \sum_{a_j \in A} e_{ij}$.

In the agent organization, each agent is also assigned a single fixed skill, $\sigma_i \in [1, \sigma]$, where σ is the number of different types of skills that are present in the organization. During the team formation process, each agent can be in one of three states: UNCOMMITTED, COMMITTED, or ACTIVE. An agent in the UNCOMMITTED state is available and not assigned to any task. An agent in the COMMITTED state has selected a task, but the full team to work on the task has not yet formed. Finally, an agent in the ACTIVE state is a member of a team that has fulfilled all of the skill requirements for a task and is actively working on that task. Only uncommitted agents can commit to a new or partially filled task. Committed agents cannot decommit from a given task. Upon task completion, agents in the active state return to the uncommitted state. We denote the state of the agent as s_i . Later we will discuss extending the team formation model to allow the agents to adapt their local network structure.

Tasks are introduced at fixed task introduction intervals, where the length of the interval between tasks is given by the model parameter, μ . Tasks are globally advertised (i.e., announced to all agents). Each task T_k has an associated size requirement, $|T_k|$, and a $|T_k|$ -dimensional vector of required skills, R_{T_k} . The skills required for a given task T_k are chosen uniformly from $[1, \sigma]$. Each task is advertised for a finite number of time steps γ , ensuring 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 successfully form to fill the requirements of a given task are only active for a finite number of time steps α . A *successful team* is one that fulfills the skill requirements for a given task within γ time steps.

The agent social network explicitly restricts the sets of agents that can form teams.

DEFINITION 3.1. A **valid team** is a set of agents $M_k = \{a_i\}$ that induce a connected subgraph of the agent social network and whose skill set $\{\sigma_i\}$ fulfills the skill requirements for a given task T_k .

The requirement of a team to induce a connected subgraph of the agent social network means that for some agent in the team, $a_i \in M_k$, there must exist at least one other agent, $a_j \in M_k, i \neq j$, such that $e_{ij} = 1$. This implies that an uncommitted agent is only eligible to commit to a task in two situations: 1) *team initiation*, when no other agents are committed to the task, and 2) *team joining*, when at least one neighbor of the agent is already committed to the task. There are many possible heuristics for initiating and joining teams, but in order to focus on network adaptation, we have selected two simple strategies.

During each iteration of the model, the agents are selected in a random order to update. Each agent in the UNCOMMITTED state in turn considers each task in a random order. If a task currently has no other agents committed to it, an agent can choose to initiate a team, and does so with a probability equal to the proportion of the agent's immediate neighbors that are currently in the UNCOMMITTED state. The probability that an agent a_i initiates a team is

$$IP_i = \frac{\sum_{a_j \in A} e_{ij} I(s_j, \text{UNCOMMITTED})}{\sum_{a_j \in A} e_{ij}}, \quad (1)$$

where $I(x, y)$ is an indicator function that returns 1 if $x = y$ and 0

otherwise. For team joining, if an agent is eligible for a team, it always joins the team. Note that agents can only be committed to, or active on, one team at a time. Figure 1 gives the pseudocode for the *JoinTeam* algorithm used for each agent. The algorithm combines team initiation and team joining. We have selected a simple team joining strategy to ensure that the benefits of network adaptation are truly a result of network adaptation. The problem of developing or learning effective team initiating and team joining policies is also important, and is included in our on-going and future work.

Algorithm 1: *JoinTeam*

input:

a_i : an agent with $s_i = \text{UNCOMMITTED}$,
 $T = \{T_1, T_2, \dots\}$: the set of current tasks,
 $M = \{M_1, M_2, \dots\}$: the set of teams, where
 M_k is associated with T_k ,
 $R = \{R_{T_1}, R_{T_2}, \dots\}$: the skill requirements for tasks in T
 E : the adjacency matrix of the agent social network

begin

for all $T_k \in T$ in random order

if $|M_k| = 0$ **and** $s_i = \text{UNCOMMITTED}$

with probability IP_i

if $\exists r \in R_{T_k} : r = \sigma_i$

$M_k \leftarrow M_k \cup \{a_i\}$

$s_i \leftarrow \text{COMMITTED}$

end if

else if $\exists a_j : e_{ij} = 1, a_j \in M_k$ **and** $s_i = \text{UNCOMMITTED}$

if $\exists r \in R_{T_k} : r = \sigma_i$ **and** r is unfilled

$M_k \leftarrow M_k \cup \{a_i\}$

$s_i \leftarrow \text{COMMITTED}$

end if

end else if

end for all

end

Figure 1: The algorithm used for each agent to decide which teams to initiate or join.

We measure the team formation performance of the agent organization as the ratio of number of teams successfully formed to the total number of tasks introduced to the system:

$$\text{org. performance} = \frac{\# \text{ of teams successfully formed}}{\# \text{ of tasks introduced}} \quad (2)$$

This measure of performance provides a global measure of how effective the agent organization is at forming teams to execute the advertised tasks. Each agents' local performance is

$$Y(a_i) = \frac{\# \text{ of successful teams joined}}{\# \text{ of teams joined}}. \quad (3)$$

In order to extend the model to include network adaptation, on each iteration, we allow agents to either attempt to join teams *or* adapt their local network structure (but not both). The details of the network adaptation strategies are presented in the next section.

4. AON STRATEGIES

In this section, we describe two strategies for AONs in the team formation environment. One strategy is based on the structural characteristics of the network and is motivated by and derived from previous finding on networked multi-agent systems [14, 29]. The

other strategy is performance-based and is intended to serve as a more realistic strategy for agent-driven network adaptation. As stated above, the goals of this paper are to determine the feasibility of agent-organized networks in a dynamic environment and to address candidate adaptation strategies and their effectiveness. Recall that we are concerned solely with network rewiring and do not allow for the explicit creation or removal of network connections.

4.1 Structure-based Adaptation Strategy

In this strategy, the agents adapt their network connectivity based on the notion of *preferential attachment* [3]. This strategy is a direct result of the finding that networks generated by the preferential attachment mechanism, so-called scale-free networks, have higher team formation performance than several other classes of network structures [14]. This strategy is intentionally simple and is intended to demonstrate that AONs are possible, at least in the context of dynamic team formation. Note that this strategy incorporates no information about performance.

The adaptation trigger for this strategy is a random process. At each iteration, an UNCOMMITTED agent chooses to adapt with a probability of $1/N$, where N is the number of agents in the system. Using this probability of adaptation results in an expected number of adaptations at each time step that does not exceed one. (The actual expected number of adaptations during any iteration is equal to the proportion of currently uncommitted agents in the system.) If an agent chooses to rewire its local connectivity, it randomly selects a current neighbor from whom to disconnect. Then the agent requests a list of neighbors from each of its neighbors and selects from among its neighbors' neighbors based on preferential attachment. Preferential attachment simply says that the probability of connecting to a given node in a network is proportional to that node's degree. Let the set of agent a_i 's neighbors' neighbors be designated as $N_i^2 = \{a_m : e_{ij} = 1, e_{jm} = 1, e_{im} = 0, m \neq i\}$. Finally, the adapting agent a_i selects an agent $a_j \in N_i^2$ with whom to establish a connection using the following probability distribution:

$$P(a_i \rightarrow a_j) = \frac{k_j}{\sum_{a_l \in N_i^2} k_l}. \quad (4)$$

The agents are prevented from establishing self-connections.

4.2 Performance-based Adaptation Strategy

Our second AON strategy is based on performance and referrals [31, 32, 34]. This strategy is intended to be more realistic and adaptive, and is based on agent performance in the organization. First, we discuss the adaptation triggering mechanism, followed by the referral process.

The adaptation trigger for the performance-based strategy has two parts. First, agents maintain a local measure of performance, $Y(a_i)$. An agent's performance measure is considered valid if the agent has attempted to join at least v teams (in the experiments described below, $v = 10$). At each iteration, agent a_i chooses to adapt its network structure if it has a valid performance measure and it is performing below the average of its immediate neighbors' performance measures (at least one of the agent's neighbors must have a valid estimate). That is, an agent chooses to adapt its network structure if

$$Y(a_i) < \frac{1}{k_i} \sum_{a_j \in A, e_{ij}=1} Y(a_j). \quad (5)$$

If an agent decides to adapt based on local performance, the rewiring is based on both performance and referrals. Specifically,

the agent will remove the connection to the immediate neighbor a_j that has the lowest valid performance:

$$a_j = \arg \min_{a_m \in A, e_{im}=1} Y(a_m). \quad (6)$$

The agent then requests a referral from its neighbor with the highest valid performance estimate. Similarly, an agent refers its neighbor with the highest valid performance estimate. Let a_l be the agent from which a_i takes the referral:

$$a_l = \arg \max_{a_m \in A, e_{ml}=1} Y(a_m). \quad (7)$$

Then the adapting agent a_i will establish a new connection with a_k , the highest-performing neighbor of a_l :

$$a_k = \arg \max_{a_m \in A, e_{ml}=1, e_{im}=0} Y(a_m). \quad (8)$$

After an agent adapts its local network structure, it resets its counters for team joining attempts and successes, in order to establish an estimate of local performance with the new network structure.

5. EXPERIMENTS AND RESULTS

We conducted a set of simulation experiments in order to evaluate the two AON strategies discussed in the previous section. In these experiments, the agent organization is initialized as an arbitrary network. We selected a modification of random geometric graphs [9] as our initial network topology. This choice was made because random geometric graphs can be constructed to have a wide range of densities (number of connections), and because their structure is based on the physical location of the agents, similar to sensor network topologies. A random geometric graph is generated by randomly placing N agents in the unit square and connecting two agents if they are within some specified distance d [9]. In order to produce networks where all agents have connectivity and to produce a range of network densities, we add a preprocessing step: after assigning random locations, we compute the minimal d that guarantees that all agents will have at least one connection. A pair of nodes is connected if the distance between the two agents is less than or equal to d . Two such spatial networks are shown in Figures 3(a) and 4(a), with their respective spring layouts in Figures 3(f) and 4(f).

The experiments we conducted started with an initial network generated with the modified random geometric graph generation process described above. Then the team formation process was run for an initial period with no adaptation (for the results presented below, 2000 iterations), so that a stable team formation performance measurement could be taken for the initial static network. After the initial period, the two AON strategies described in the previous section were enabled and the networks were allowed to evolve according to these strategies. It is important to note that at any iteration, an agent can *either* adapt their network *or* attempt to join teams, but not both: the two functions are mutually exclusive. In order to assess the utility of the strategies, we recorded the relative performance during adaptation as compared to the measurement at the end of the initial period. In these experiments, organizational performance is the percentage of tasks for which teams successfully form.

In all of the experiments presented in this paper, the parameters of the team formation model were $N = 100$, $\alpha = \gamma = \sigma = |T| = 10$, and $\mu = 2$. Given these parameter settings, the network is heavily loaded with tasks. It is important to note that this is a sampling of the results; similar results were found for a range of team formation parameters. All results are the average of 50 simulations with a different initial network for each simulation.

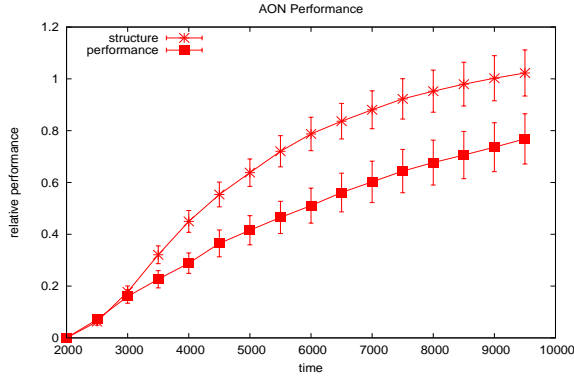


Figure 2: Performance change over time as a result of local network adaptation.

Figure 2 shows the results for relative team formation performance the AON strategies. The error bars represent one standard deviation from the mean. A value of 0 on the y-axis represents no change in performance; a value of 1 represents a 100% increase in performance. The structure-based adaptation strategy (labeled “structure” in the figures) nearly doubles the organizational performance on average after 7,500 iterations with network adaptation, with the performance-based strategy not far behind. This result demonstrates the practicality and feasibility of AONs for dynamic team formation. It was expected that the structure-based adaptation strategy would result in dramatically increased performance, since it is directly searching for a network topology that is known to outperform other network structures in highly loaded team formation processes (i.e., a network with short average path lengths between the nodes and a hub structure) [14]. The dramatic increase in performance as a result of the performance-based strategy is more surprising. An interpretation as to why this increase occurs will be discussed below. Note that for both strategies, the rate at which performance is improving slows over time, suggesting the possibility for the networks to stabilize if the adaptation triggers are adjusted to account for changes in performance over time.

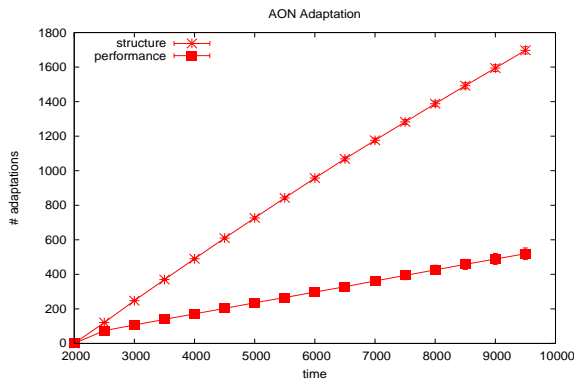


Figure 3: Cumulative number of adaptations over time.

We also monitored the number of adaptations for each of the AON strategies over time. Figure 3 shows the cumulative number of adaptations for the two strategies. Notice that the performance-based adaptation strategy performed significantly fewer adaptations than the structure-based strategy. Qualitatively, the space between

the performance curves is noticeably smaller than the space between the adaptation curves. This implies that the performance-based strategy may be more efficient (i.e., each adaptation has a larger impact on performance). This would be especially important in environments where network adaptation is expensive from a cost or resource perspective.

In order to understand how the two strategies produce network structures with increased team formation performance, we randomly selected two instances for which to show the network structures as they evolved. Figure 4 shows the network topology of the AON as it evolves with the structure-based strategy. The top row shows the spatial layout of the network given the agent positions; the bottom row shows the logical network topology, drawn with a spring layout. As expected, the structure-based adaptation strategy drives the global network topology toward a network with a hub structure and short average path length in each of the components.¹ This is evidenced qualitatively by the compact layout of the logical structures (bottom row of the figure). The hub structure can also be observed directly from the geographic layout. Because the adaptation strategy is designed for a specific graph topology, the adaptation toward the hub-structured network is rapid (apparent at $t = 4000$).

Figure 5 shows the evolution of a network under the performance-based adaptation strategy. It should be noted that the starting networks presented in Figures 4 and 5 are different. Unlike the structure-based strategy, the performance-based strategy approaches a network structure with short average path length more slowly. This is evidenced by the gradual compacting of the largest cluster over time in the bottom row of the figure. The fact that the strategy leads to a network structure with short average path length is intuitive when considering the mechanisms of the strategy. The strategy essentially has agents rewiring from lower-performing agents to higher-performing agents. Agents that have higher performance over time tend to be agents that have high degree, since there are many more opportunities to join teams and many more possible teams to form in their local neighborhood.

In summary, both adaptation strategies led to more efficient organizational structures. Interestingly, the performance-based adaptation strategy leads to a qualitatively similar network structure as the structure-based strategy – that of a network with short average path length and a hub structure.

6. RELATED WORK

In the multi-agent systems community, the ideas of network adaptation and learning have been studied in the context of referral networks [31, 32, 34]. In these studies, agents are producers and consumers of services (to include information). Consumer agents request services of the agents they are connected to in a social network. The neighboring agents can take one of three subsequent actions: 1) no response, 2) provide the service, or 3) offer a referral. Learning occurs through agents building models, including models of expertise and sociability, of neighboring and more distant agents in the system. Based on these models, agents modify their local social network structure with the goal of discovering and providing services more effectively. Outcomes of this work include an analysis of the emergent community structure with suggestions for methods of discovering community structure [32] and strategies for effectively searching the network for service providers [33]. We have incorporated the ideas from the study of referral networks in

¹In both of the adaptation strategies, once agents are disconnected, they can never be reconnected, because there is no way for the other agents to “know” they are out there. Strategies that incorporate reconnections are left as future work.

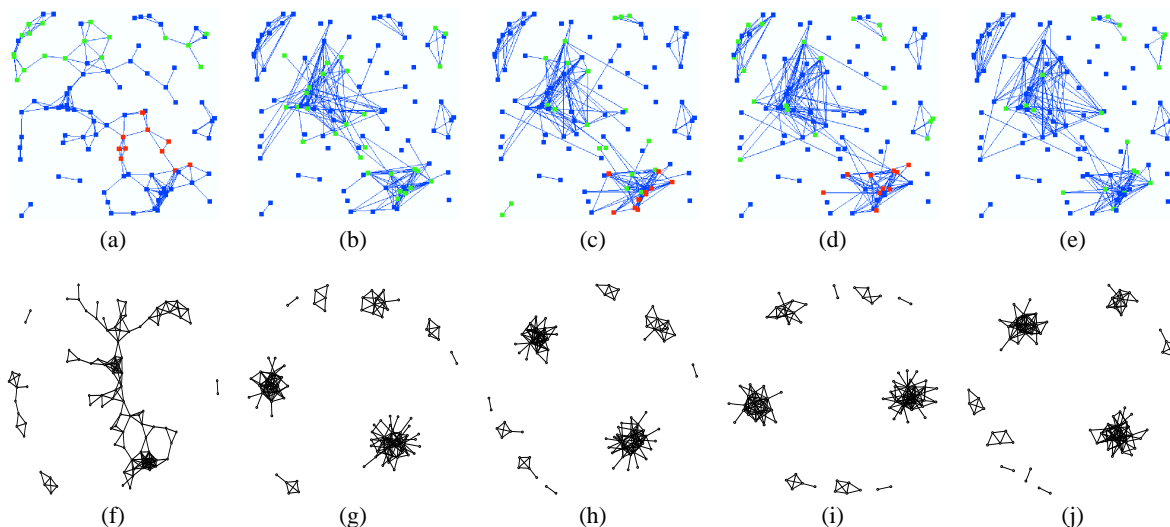


Figure 4: Evolution of the network under the structure-based adaptation strategy. The top row is the spatial layout with the agents in their initial positions. The bottom row is a non-spatial spring layout. The positions of the clusters in the latter is arbitrary. The five samples from left to right are: (a) and (f) $t = 2000$; (b) and (g) $t = 4000$; (c) and (h) $t = 6000$; (d) and (i) $t = 8000$; and (e) and (j) $t = 10000$. The network has 217 connections and the performance was increased from 0.0423 at $t = 2000$ to 0.0675 at $t = 10000$. (The colors of the nodes in the spatial layout, blue, green, and red, correspond to UNCOMMITTED, COMMITTED, and ACTIVE, respectively.)

producer-consumer environments and applied them to the problem of team formation. Additionally, we have removed the necessity to model other agents in the organization and have shown that without modeling, effective network structures can still organize. Strategies that allow agents to build models of other agents in the organization may prove useful in the future.

In other fields, network dynamics has emerged as an important topic of study. Since many physical systems have a network interpretation (i.e., system of interdependent parts), physicists have developed techniques for modeling and analyzing networks that change structure over time. In one such study, reinforcement learning is used as a mechanism for discovering which agents to follow in the minority game [4]. Other work focuses on the stability of dynamic processes on networks, self-assembly, and models that replicate properties of networks found in nature [11, 25]. Finally, in a study at the intersection of social science and statistics, a simple model of a “stochastically evolving social network” is used to guide the pairings for repeated games [26].

7. ON-GOING AND FUTURE WORK

While we have demonstrated that agent-organized networks are feasible for improving collective performance for team formation, there is still much to be done. Currently, we are pursuing quantitative analyses of resulting network structures for various adaptation strategies. It will also be important to analyze AON strategies that both increase and decrease global performance. Additionally, we are exploring other networked multi-agent system domains to determine if AON strategies can apply to broad classes of systems. These systems include market environments, supply networks, and sensor networks. Our long term goal is to develop a principled

framework for designing and employing agent-organized network mechanisms.

In related work, we are also exploring methods for learning team joining policies and for adapting individual agent skills to improve performance. These other organizational learning paradigms share some of the same challenges as organizational learning through AONs.

8. CONCLUSION

We have introduced the notion of agent-organized networks and shown that they are a feasible mechanism for distributed, on-line learning of effective organizational structures in a dynamic team formation environment. Through the design and application of two different strategies, we have shown that very different mechanisms can lead to efficient network structures. The important implication of agent-organized networks is that efficient organizational structures can emerge as a result of individual agent decisions based solely on local, and perhaps imperfect, information. Additionally, AONs provide a mechanism for real-time organizational learning, embedding the adaptation as part of the function of the agents in the organization.

9. ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers and Blazej Bulka for insightful and helpful comments for improving this paper.

10. REFERENCES

- [1] S. Abdallah and V. Lesser. Organization-based cooperative coalition formation. In *Proceedings of the IEEE/WIC/ACM*




Figure 5: Evolution of the network under the performance-based adaptation strategy. The top row is the spatial layout with the agents in their initial positions. The bottom row is a non-spatial spring layout. The positions of the clusters in the latter is arbitrary. The five samples from left to right are: (a) and (f) $t = 2000$; (b) and (g) $t = 4000$; (c) and (h) $t = 6000$; (d) and (i) $t = 8000$; and (e) and (j) $t = 10000$. The network has 294 connections and the performance was increased from 0.0697 at $t = 2000$ to 0.142 at $t = 10000$. (The colors of the nodes in the spatial layout, blue, green, and red, correspond to UNCOMMITTED, COMMITTED, and ACTIVE, respectively.)

International Conference on Intelligent Agent Technology (IAT), September 2004.

- [2] G. Abramson and M. Kuperman. Social games in a social network. *Phys. Rev. E*, 63(030901), 2001.
- [3] R. Albert and A. Barabási. Statistical mechanics of complex networks. *Review of Modern Physics*, 99(3):7314–7316, May 2002.
- [4] M. Anghel, Z. Toroczkai, K. Bassler, and G. Korniss. Competition-driven network dynamics: Emergence of a scale-free leadership structure and collective efficiency. *Physical Review Letters*, 92(5), 2004.
- [5] R. C. Arkin and J. Diaz. Line-of-sight constrained exploration for reactive multiagent robotic teams. In *Proceedings of the Seventh International Workshop on Advanced Motion Control (AMC'02)*, 2002.
- [6] R. Arunachalam, N. Sadeh, J. Eriksson, N. Finne, and S. Janson. The supply chain management game for the Trading Agent Competition 2004. Technical Report CMU-CS-04-107, Carnegie Mellon University School of Computer Science, July 2004.
- [7] R. Axtell. Effects of interaction topology and activation regime in several multi-agent systems. In *MABS*, pages 33–48, 2000.
- [8] D. Culler, D. Estrin, and M. Srivastava. Overview of sensor networks. *IEEE Computer*, 37(8):41–19, 2004.
- [9] J. Dall and M. Christensen. Random geometric graphs. *Phys. Rev. E*, 66(016121), 2002.
- [10] J. Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141:171–185, 2002.
- [11] H. Ebel, J. Davidsen, and S. Bornholdt. Dynamics of social networks. *Complexity*, 8(2):24–27, 2002.
- [12] M. Fox, M. Barbuceanu, and R. Teigen. Agent-oriented supply-chain management. *International Journal of Flexible Manufacturing Systems*, 12(2/3):165–188, 2000.
- [13] M. Gaston and M. desJardins. Team formation in complex networks. In *Proceedings of the 1st NAACSOS Conference*, June 2003.
- [14] M. Gaston, J. Simmons, and M. desJardins. Adapting network structures for efficient team formation. In *Proceedings of AAMAS-04 Workshop on Learning and Evolution in Agent-based Systems*, July 2004.
- [15] P. Holme, A. Trusina, B. J. Kim, and P. Minnhagen. Prisoners' dilemma in real-world acquaintance networks: Spikes and quasi-equilibria induced by the interplay between structure and dynamics. *Phys. Rev. E*, 68(030901), 2003.
- [16] B. Horling, B. Benyo, and V. Lesser. Using self diagnosis to adapt organizational structure. In *5th International Conference on Autonomous Agents*, 2001.
- [17] L. Hunsberger and B. Grosz. A combinatorial auction for collaborative planning. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS-2000)*, 2000.
- [18] B. Kim, A. Trusina, P. Holme, P. Minnhagen, J. Chung, and M. Choi. Dynamics instabilities induced by asymmetric influence: Prisoners' dilemma game in small-world networks. *Phys. Rev. E*, 66(021907), 2002.
- [19] S. Kraus, O. Shehory, and G. Taase. Coalition formation with uncertain heterogeneous information. In *Proceedings of the*

Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '03), July 2003.

- [20] P. J. Modi, H. Jung, M. Tambe, W.-M. Shen, and S. Kulkarni. A dynamic distributed constraint satisfaction approach to resource allocation. In *Proceedings of the Seventh International Conference on Principles and Practice of Constraint Programming*, 2001.
- [21] R. Nair, M. Tambe, and S. Marsella. Team formation for reformation. In *Proceedings of the AAAI Spring Symposium on Intelligent Distributed and Embedded Systems*, 2002.
- [22] R. Nair, M. Tambe, and S. Marsella. Role allocation and reallocation in multiagent teams: Towards a practical analysis. In *Proceedings of the Second International Joint Conference on Agents and Multiagent Systems (AAMAS)*, 2003.
- [23] M. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- [24] D. Pynadath and M. Tambe. Multiagent teamwork: Analyzing the optimality and complexity of key theories and models. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS '02)*, 2002.
- [25] F. Schweitzer and B. Tilch. Self-assembling of networks in an agent-based model. *Physical Review E*, 66:1–9, 2002.
- [26] B. Skyrms and R. Pemantle. A dynamic model of social network formation. *Proceedings of the National Academy of Sciences USA*, 97:9340–9346, 2000.
- [27] A. Szolnoki and G. Szabo. Phase transitions for rock-scissors-paper game on different networks. To appear in *Phys. Rev. E*, 2004.
- [28] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.
- [29] H. P. Thadakamalla, U. N. Raghaven, S. Kumera, and R. Albert. Survivability of multiagent-based supply networks: A topological perspective. *IEEE Intelligent Systems*, 19(5):24–31, 2004.
- [30] D. Watts and S. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.
- [31] P. Yolum and M. P. Singh. Dynamic communities in referral networks. *Web Intelligence and Agent Systems (WIAS)*, 1(2):105–116, 2003.
- [32] P. Yolum and M. P. Singh. Emergent personalized communities in referral networks. In *Proceedings of the IJCAI Workshop on Intelligent Techniques for Web Personalization (ITWP)*, August 2003.
- [33] B. Yu and M. P. Singh. Searching social networks. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. ACM Press, July 2003.
- [34] B. Yu, M. Venkatraman, and M. P. Singh. An adaptive social network for information access: Architecture and experimental results. *Applied Artificial Intelligence*, 17(1):21–38, January 2003.