Agent-Based Sensor Coalition Formation

Robin Glinton, Paul Scerri, Katia Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, U.S.A.
Email: rglinton,pscerri,katia@cs.cmu.edu

***Abstract*—Large numbers of heterogeneous sensors the collect and fuse information in dynamic environments are envisioned in domains, such as military operations, disaster response or border surveillance. Depending on the dynamically evolving needs, operators will submit to groups of these sensors information acquisition and fusion goals that must be fulfilled within time constraints. To fulfill these goals, the sensors must autonomously and dynamically form coalitions. The problem that this paper addresses is the dynamic coalition formation problem and the evolving performance of coalitions over time. Coalition formation is an NP complete problem. One of the ways to mitigate the computational cost is to constrain the decentralized coalition formation problem by taking into consideration the underlying network structure. [1] has showed (a) that the network topology has a significant effect on the quality of the formed coalitions and their performance, and (b) that it is possible to develop agents that intelligently adapt the network structure to increase the ability of the organization to form good quality coalitions. However, that work used unrealistically simple coalitional models and did not perform an analysis of the network topologies and adaptation policies that could result. In this paper we make three contributions. First we present an analysis and results on the underlying network topologies that are formed. Second, we develop and analyze a more realistic coalition formation model. Third we present two new network adaptation policies. Experimental evaluation of our contributions are presented.**

I. Introduction

In an increasing range of applications, large numbers of heterogeneous sensors will be used to collect and fuse data in dynamic, partly observable environments. In domains, such as military operations, disaster response or border surveillance, groups of these sensors will be dynamically tasked to provide information about some feature of the environment. When these teams are large, it will be infeasible for operators to individually tasks assets, instead high level information acquisition goals will be submitted to the team and coalitions of sensors dynamically formed to provide the information. This paper addresses the issue of the formation of these coalitions/groups and their performance over time.

Coalition formation is an extensively studied problem, in microeconomics, game theory and computer science, in particular in multi-agent systems [2], [3], [4], [5]. A coalition is a set of agents (considered self-interested in the classical formulation) that cooperate to achieve a common goal. The classical notion of coalition formation is for the agents to form a coalition that maximizes their own individual payoff while at the same time maximizes group (joint) payoff and guarantees stability. A coalition is stable if the payoff obtained by the agents are such that both individual and group rationality are satisfied. In other words, the agents are seeking to form a coalition in such a way that the payoffs to the agents are such that no individual (or subgroup of) agent(s) has an incentive to depart from the coalition. In our work, the agents are not self interested, their utility arises from their desire to provide the best contribution to the group utility. The coalition formation problem is NP-complete [6] and most of the work in the literature deals with the centralized and static problem[7]. Work in decentralized coalition formation [5], [8] provide schemes mainly focused on optimizing the group's total utility. Most of this previous work restricts the size of the coalition in order to reach a computationally efficient solution. Gaston [1] is one of the first researchers to consider the imposition of network structure on the decentralized coalition formation problem in order to (a) mitigate high computational cost, and (b) mitigate coalition instability by assuming that only agents currently connected to the coalition initiator can be candidate coalition members. Additionally, Gaston's work considers the behavior of coalitions over time as tasks to be solved by coalitions arrive dynamically and as the agents adapt the network structure in order to better address the performance of these tasks. Gaston showed that if the network is allowed to change intelligently over time, the overall group can still form high quality coalitions[1]. However, Gaston neither investigated the nature of the networks that the group eventually created, nor did he use a realistic model of coalitions. Thus, while Gaston's results are impressive, their practical applicability to real world sensor coalition problems was not established.

The first contribution of this paper is an examination of the networks formed by Gaston's algorithms for network adaptation. It was found that, in nearly all cases, scale-free networks, characterized by a small number of "hub" nodes with very many links, resulted after the network was able to adapt. These hub nodes explain the good overall performance after adaptation – since the hub nodes are connected to many other agents, they provide a locus around which many coalitions can be formed. However, scale-free networks are not always desirable, especially in the case of sensors, since a high communication and computational load is placed on the hub nodes. Therefore an examination of the network adaptation algorithm to determine coalition behavior over time was undertaken while limiting the maximum connectivity of nodes. It was found that no substantial drop-off in performance was observed, showing that the network could be successfully adapted to allow high quality coalitions to be formed, without the undesirable high degree of connectivity of hub nodes.

The second contribution of this work is to introduce a more realistic model of coalitions. Specifically, a model is presented where the value of the coalition is a non-linear function of the scalar capabilities of the members in the coalition. In the sensor network domain, for example, some assets are better than others perhaps because of higher fidelity cameras; some assets may provide higher group utility by complementing each other, perhaps by providing information that allows the other to better focus its resources; on the other hand, some sets of sensor assets may provide lower utility to the group since they could interfere with each other, perhaps by producing signals that cause noise for another sensor. For this more realistic model, the purely structure-based network adaptation process used by Gaston does not provide any performance benefits.

The third contribution of this paper is two new adaptation policies that work with the more realistic coalition model. One policy encourages agents to connect to other agents with distinctly different capabilities. The second policy encourage agents to make connections to agents that have previously been performing well or performing poorly to try to improve overall performance.

Extensive simulations were performed, modeling a wide range of large heterogeneous sensor teams. The key things we found were that a policy which rewires local network connectivity based solely on network topology has a negligible effect on coalition formation performance when a realistic coalition model is used. This is in stark contrast to the significant performance improvement that results when a simplistic coalition model is used. In addition we found that a policy that preferentially attatches agents to those with higher performance produces a significant improvement in performance. Conversely a policy which that preferentially attatches agents to those with dissimilar capabilities actually significantly decreases performance.

II. Problem Statement

The following is a formal description of the problem addressed in this work. There is a population of $|A|$ agents represented by the set $A = \{a_1, \ldots, a_{|A|}\}$. The agents are connected by a network modeled by a matrix E with elements e_{ij}, where $e_{ij} = 1$ indicates that an edge exists between a_i and a_j. Each agent is assigned a vector of capabilities $C = \{c_1, \ldots, c_{|C|}\}$ with elements taken from the range $c_i \in R$ where $c_i \in [0, 1]$.

Agents must form coalitions given by $M_k \subset A$ on connected sub-graphs, to complete tasks introduced to the population and $G^{[t_1,t_2]}$ denotes the set of successful coalitions formed during the interval $[t_1, t_2]$. Tasks are globally advertised and introduced at fixed intervals μ. Each task T_k has a size requirement $|T_k|$ and a vector of required capabilities $R_{T_k} = \{r_1, \ldots, r_{|R_{T_k}|}\}$, where $r_i = 1$ if task T_k requires capability c_i and $r_i = 0$ otherwise. Tasks are advertised for γ time steps. Any agents committed to a task are freed if the full compliment of capabilities does not become available in the time window during which the task is advertised. Tasks take α time steps to complete. For agent a_j to join a coalition M_k there must exist an edge $e_{ij} = 1$ such that $a_i \in M_k$. Each coalition M_k has a value V_k where:

$$V_k = \prod_i \max_{a_j \in M_k} c_i^j$$

The objective is to optimize the following function:

$$\max_E \sum_{t=0}^{\infty} \sum_k V_k$$

III. Coalition Formation

This section describes the algorithm used to form coalitions on a network. The pseudocode for this algorithm is given by Algorithm 1.

During execution of Algorithm 1, Agents can be in one of three states, *active*, *committed*, OR, *uncommitted*. An agent in the *active* state is a member of a coalition which is in the process of executing a task. An agent in the *uncommitted* state has not joined a coalition and an agent in the *committed* state is a member of a coalition which lacks the full compliment of capabilities needed to accomplish its associated task.

At each time step, each agent in the *uncommitted* state has the opportunity to adapt its local connectivity in the network by breaking a link with one of its neighboring agents and forming a new link with another. Each agent chooses to adapt connectivity with probability $1/|A|$ where $|A|$ is the total number of agents. The various policies used by agents to adapt local network connectivity are described in Section V.

If an agent in the *uncommitted* state does not adapt its local connectivity, it attempts to form a new coalition or join an existing one using Algorithm 1. An agent considers each task T_k in turn. If the agent encounters a task T_k with no committed agents, $|M_k| = 0$, the agent attempts to initiate a coalition to address T_k with probability P_I, (lines 4-7 of Algorithm 1).

The probability P_I that an agent will initiate a coalition to address a task T_k is proportional to the number of its immediate neighbors in the *uncommitted* state and the capabilities that the agent has that are needed to fulfill the task. Formally:

$$P_I = \frac{\prod_{l:r_l^k=1} c_l^j \times \sum_{a_j \in A} e_{ij} I(s_i, \mathit{uncommitted})}{\sum_{a_j \in A} e_{ij}}$$

where $I(x, y) = 1$ when $x = y$ and 0 otherwise.

If an agent encounters a task with committed agents (a coalition has already been initiated to address the task) it can attempt to join the coalition if it meets two requirements. It must be connected by a link to an agent already in that coalition and must also have a capability required by that coalition. If these requirements are met the agent attempts to join the coalition with probability P_j (lines 8-13 of Algorithm 1).

The probability P_j that an agent a_j will join a coalition that has already been initiated to fulfill a task T_k is proportional to the capabilities of the agent required by the task:

$$P_j = \prod_{l:r_l^k=1} c_l^j$$

Algorithm 1: Algorithm used by an agent to initiate or join a coalition.

JOINCOALITION()
foreach $T_k \in T$ in random order
    if $|M_k| = 0$ and $s_i =$ *uncommitted*
        $diceRoll \leftarrow$ UNIFORMRANDOM($[0, 1]$)
        if $diceRoll < P_I$
            if $\exists r \in R_{T_k}$: r is unfilled
                $M_k \leftarrow M_k \bigcup \{a_i\}$
                $s_i \leftarrow$ *committed*
    else if $\exists a_j : e_{ij} = 1, a_j \in M_k$ and $si =$ *uncommitted*
        $diceRoll \leftarrow$ UNIFORMRANDOM($[0, 1]$)
        if $diceRoll < P_J$
            if $\exists r \in R_{T_k}$: r is unfilled
                $M_k \leftarrow M_k \bigcup \{a_i\}$
                $si \leftarrow$ *committed*

IV. STRUCTURE BASED POLICY

Given the above coalition formation algorithm, the underlying network structure is clearly critical. The primary contribution of Gaston's work was a network adaptation algorithm that improves the orginizations ability to form coalitions. We briefly review that algorithm here, before analyzing the networks that result.

When an agent elects to adapt its local connectivity, it probabilistically selects from its neighbor's neighbors for candidates to form links to biased by the network degree of the candidates and it will choose uniformly at random from its existing links for a link to drop. Agent $a_i's$ neighbor's neighbors are given by $N_i^2 = \{a_m : e_{ij} = 1, e_{jm} = 1, e_{im} = 0, m \neq i\}$. When a_i adapts, it selects an agent $a_j \in N_i^2$ to establish a link to using the following probability distribution:

$$P(a_i \to a_j) = \frac{\text{number of links } a_j}{\sum_{a_l \in N_i^2} \text{number of links } a_l}$$

The results reported in [1] showed that a simple algorithm could be used to adapt a wide variety of initial network topologies producing a network structure on average 100% more efficient than the starting network.

A. Policy Analysis

Our first objective in this investigation was to understand both the structure of the networks formed by the structure based policy used in [1] and the relationship between the resulting network structure and the rate of task completion, and hence value, of coalitions formed on the network.

The model of coalition formation used in Gaston's work differs from the one used in this paper as follows. In Gaston's work each agent is assigned a single capability given by $\sigma_i \in [1, \sigma]$ where σ is the number of available capabilities.

Each task T_k has a size requirement $|T_k|$ and a vector of required capabilities R_{T_k} of size $|T_k|$. Capabilities are chosen uniformly from $[1, \sigma]$.

$$P_I = \frac{\sum_{a_j \in A} e_{ij} I(s_i, \mathit{UNCOMMITTED})}{\sum_{a_j \in A} e_{ij}}$$

In this model of coalition formation, an agent will always join a pre-existing coalition if it has the required capability and it is linked to another agenthat is already a member of the coalition.

The performance metric used in [1] is given by:

$$\text{Performance}_t = \frac{\text{tasks completed on the interval } [t - 1000, t]}{\text{tasks completed on the interval } [0, 2000]}$$

where time is measured in discrete dimensionless units.

For our initial network topology, following Gaston, we constructed a variation of the random geometric graph. This is accomplished by randomly distributing $|A|$ points, one for each agent, in the unit square. Links are made between agents whose corresponding points have a euclidean distance between them that is less than d.

Each graph presented through the remainder of the paper shows data points that are averaged over 50 trials for a population of 500 agents. The following parameter values are common to all trials: $\alpha = \gamma = \sigma = |T| = 10$, $\mu = 2$.

For each trial the agents were allowed to complete tasks without network adaptation for 2000 iterations to establish a task completion baseline. Past 2000 iterations network adaptation was turned on.

Figure 1 shows a histogram of the links per agent, for the resulting network topology after the structure based network adaptation algorithm is run for 28000 iterations. The large number of agents with relatively few links trailing off to a very small number of agents with relatively large degree is a signature of a scale free network. The fact that the local broadcast algorithm is forming scale free networks explains the high performance of the algorithm. The small number of agents with high degree connect most of the other agents together making it possible to form connected subgraphs with all required capabilities to complete a task with high probability. This is fine with a model like the one discussed in this section where coalition value is solely dependent on the presence or absence of a certain capability. However, in coalition models which more accurately reflect the synergies which exist between groups of real sensors, capabilities are non-binary. The value of coalitions formed are then dependent on the magnitude of a particular capability within a coalition. This in contrast to the simple model presented in this section where the value of a coalition is dependent only on the presence or abscensce of a capability. When the value of a coalition is dependent on more than just the presence of

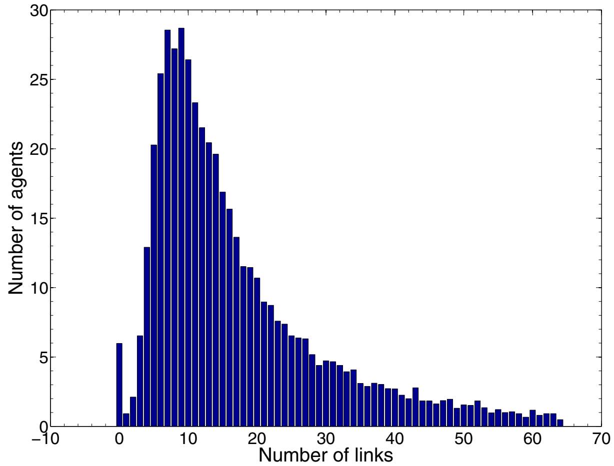

Figure 1. The histogram, of links per agent, resulting from the use of the structure based policy and simple coaltion model (binary valued capabilities).

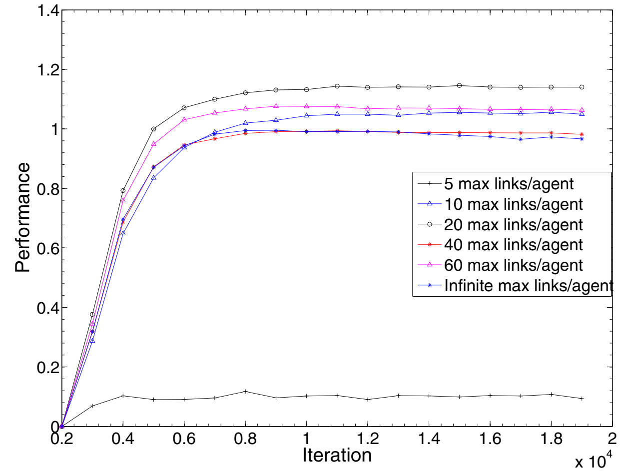

Figure 2. Effect of limiting the max links/agent on performance of structure based adaptation.

required capabilities, forming a scale free network without considering the magnitude of the capabilities of agents is not likely to yield high coalition value.

Furthermore, From an effectiveness standpoint, scale-free networks are fine, however, there are various reasons to prefer not to have scale-free networks. This undesirability stems from the concentration of links and hence communication and effort at the hub. Thus it is desirable to see whether these hubs are necessary. Figure 2 shows the effect of limiting the maximum allowable number of links per agent on the local broadcast algorithm. The graph shows that a maximum number of links greater than 10 has no impact on the network adaptation performance of the local broadcast algorithm. However, below 10 maximum links per agent, performance drops significantly. This is because there were a total of 10 possible capabilities uniformly distributed within the agent population. With a population of 500 agents this gives an expected number of 50 agents with each capability. This means that any degree greater than 10 gives an agent forming a coalition a high probability of having links to all of the necessary capabilities. This is important because it means that the hubs with high degree formed by the local broadcast algorithm waste bandwidth by communicating along many unecessary links.

V. Token Algorithm for Network Adaptation

When an agent is not currently a member of a coalition it will attempt to change its neighbors in the network in an attempt to produce a new network topology that results in higher organizational performance. This section describes an algorithm used by an agent to locate a potential candidate to create a new link with. Two separate policies, the performance based and similarity based policies, are then used by the agent to decide whether to actually form the new link. These policies are described in Sections V-A and V-B respectively.

We chose to use token algorithms because token algorithms have been shown to collect large amounts of information while using relatively small amounts of bandwidth [9].

When an agent a_i decides to form a new link it will break an existing link by removing a link to a neighbor, this dropped neighbor is selected relative to the policy the agent is using. Details of the selection process will be described in Sections V-A and V-B which describe the two policies. Next the agent will instantiate a token. This token contains a record of the best agent to rewire to, of those visited by the token, relative to a policy specific function $Q(a_i, a_j)$ which gives a metric for the quality of the match between agents a_i and a_j for belonging to the same coalition and hence rewiring. The initiating agent a_i passes the token to a neighbor using the PassBiased(token) method given by Algorithm 3, which is used by all algorithms presented in this paper to pass tokens between agents.

Upon receiving a token an agent a_j will use Algorithm 2 to process the token. When an agent receives a token it first checks if it is the first recipient of the token and passes it on if this is the case (lines 3-5). This is because if an agent is the first recipient of a token it is already linked to the agent a_i that instantiated the token. Otherwise, if the TTL of the token is non-zero, the agent will compare itself to the best agent a_{best} visited by the token thus far relative to Q. If $Q(a_i, a_j) > Q(a_i, a_{best})$ then agent a_j will update the token to reflect that $a_{best} = a_j$ and will pass it on (lines 8-11). If the TTL of the token is zero then the token is returned to the initiating agent. The initiating agent will then form a link with the agent, a_{best} with the best performance of those visited by the token.

The algorithm passBiased(token) used in lines 4 and 12 of Algorithm 2 is given by Algorithm 3. This algorithm is used by an agent to select a neighbor to pass a token to biased by the network degree of its neighbors.

Algorithm 2:

```
FINDTEAMMATETOKEN()
foreach t ∈ tokenList
    t.TTL ← t.TTL − 1
    if A_visited = ∅
        A_visited.add(this)
        PASSBIASED(t)
    else if t.TTL = 0
        CREATEEDGE(t.a_i, t.a_best)
    else if Q(t.a_i, t.a_best) < Q(t.a_i, t.a_j)
        t.a_best = t.a_j
        PASSBIASED(t)
    else
        PASSBIASED(t)

PASSBIASED(token)
totalDegree ← 0
while true
    n ← GETRANDOMNEIGHBORUNIFORM()
    totalDegree ← totalDegree + n.degree
    r ← UNIFORMRANDOM([0, 1])
    if r < n.degree / totalDegree
        PASSTOKEN(token, n)
        break
```

A. Performance Based Policy

The first policy used by agents to adapt local network connectivity is based on the performance of individual agents. Each agent maintains a record of its performance since the last time it adapted its own local connectivity. An agent will reset its performance data after adaptation. Performance data is considered valid after an agent has attempted to join at least 10 coalitions.

When an agent decides to form a new link it will break an existing link by removing a link to its neighbor that has valid performance data and the lowest performance.

For this policy:

$$Q(a_i, a_j) = \begin{cases} 0 & p_{a_i} < p_{a_j}, \\ p_{a_i} - p_{a_j} & p_{a_i} > p_{a_j} \end{cases}$$

where, for $a_i \in M_k$, $p_{a_i} = \sum_{M_k \in G^{[t_m, t_n]}} V_k$ and t_m,t_n are the last time a_i rewired and the current time step respectively.

We ran an experiment to test the network adaptation performance of the Similarity Based Policy. The performance metric used throughout this paper when the complex coalition model is employed is given by:

$$\text{Performance}_t = \frac{\sum_{M_i \in G^{[t-1000,t]}} V_i}{\sum_{M_j \in G^{[0,2000]}} V_j}$$

where time t is measured in discrete dimensionless units.

Figure 3 shows the result of the experiment. The figure shows that organizational performance improved by about 50% using the performance based policy. The good performance is a result of the fact that better performing agents tend to have higher capabilities and the performance based policy encourages these more capable agents to have a higher network degree and hence be a member of more coalitions which in turn increases coalition value.

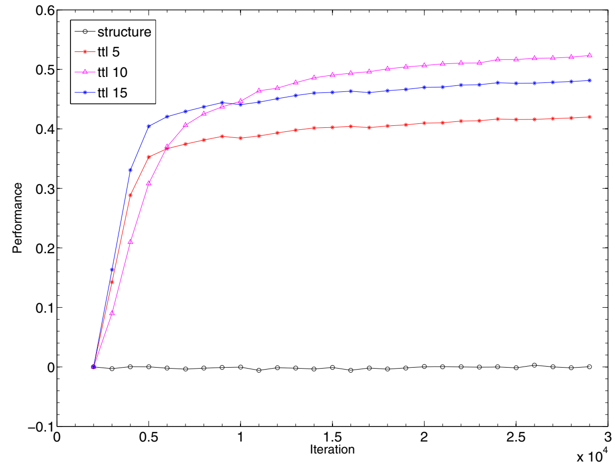

Figure 3. The performance based network adaptation policy results in up to a 50% improvement in organization performance. However, the structure based policy has no effect on performance.

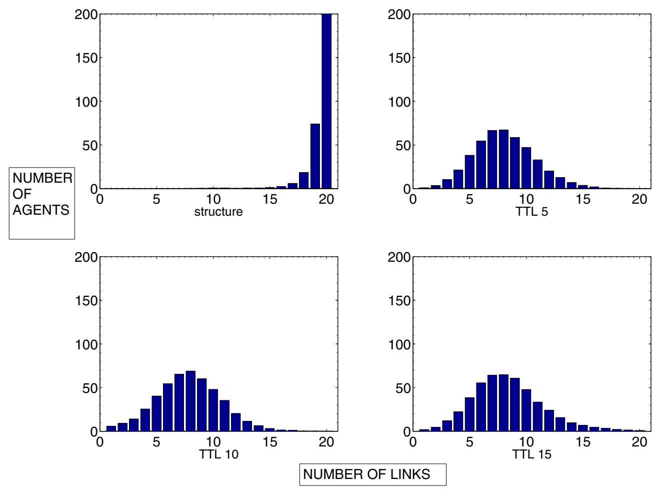

Figure 4. Histogram of the resulting network topology, links per agent, after 28000 iterations of the structure based policy and the performance based policy for TTLs of 5,10, and 15.

B. Similarity Based Policy

The second rewiring policy used by agents to adapt local network connectivity is based on the similarity between the capabilities of individual agents. This policy is meant to encourage coalitions of agents with diverse capabilities. This is based on the idea that if agents are more likely to have neighbors with complimentary capabilities, then the value of coalitions formed will be higher on average since the value of a coalition is a function of the maximum of each capability available in the coalition.

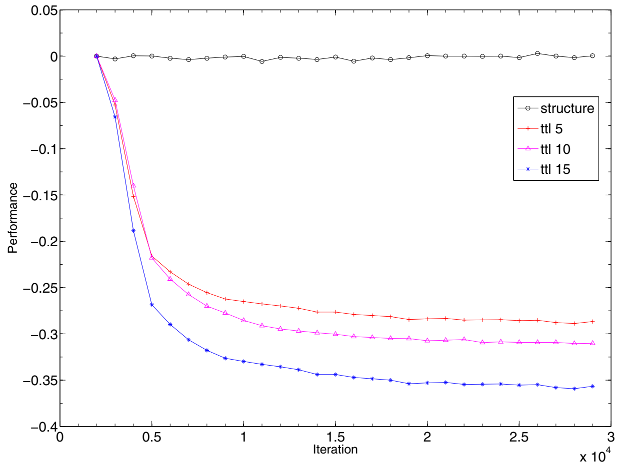

Figure 5. The similarity based network adaptation policy results in a 36% reduction in organization performance. However, the structure based policy has no effect on performance.

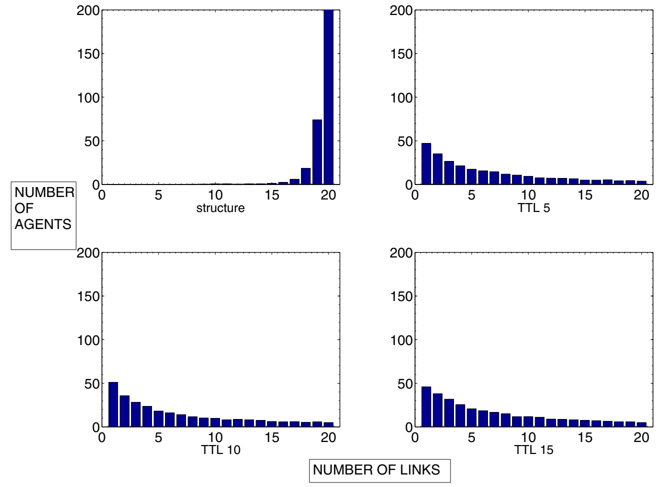

Figure 6. Histogram of the resulting network topology, links per agent, after 28000 iterations of the structure based policy and the similarity based policy for TTLs of 5,10, and 15.

The measure of similarity used in this paper is the cosine similarity between the vectors of agent capabilities given below.

$$\mathrm{Sim}(a_i, a_j) = \frac{\sum_l c_l^i \times c_l^j}{||C_i||||C_j||}$$

For the Similarity based policy $Q(a_i, a_j) = 1 - Sim(a_i, a_j)$. That is, based on this policy, the probability of two agents having a link, and hence the probability of joining the same coalition, increases with the dissimilarity of their capability vectors.

We ran an experiment to test the network adaptation performance of the Similarity Based Policy. Figure 5 shows the result of the experiment. The figure shows that after about 13000 iterations of the token algorithm using the Similarity based policy, organizational performance has actually dropped by about 35% as a result of network adaptation.

Insight into these surprising results can be found through study of Figure 6 which gives the histogram of links/node of the network that results from the operation of the Similarity based policy. Initially the links per node are normally distributed with a mean of about 7 links per node. Figure 6 shows that the Similarity based policy has drastically altered the network topology with the majority of nodes having only a single link, and the smallest number of nodes having the maximum of 20 nodes per link. This explains the poor coalition performance of the network formed. Because many of the nodes have a number of links that is significantly less than the average number of capabilities x required to form a coalition, such a structure will form significantly fewer coalitions ththe initial network structure which has an average number of links that is equal to the average number of capabilities required to form a coalition.

VI. Related Work

Many authors have looked at coalition formation in multi-agent systems [10], [11], [8]. While some work, like the Gaston work built on here, focuses on protocols for efficient coalition formation [1], other work looks at properties of the coalitions formed or puts bounds on algorithm performance. For example, Sandholm shows the properties of coalitions formed when agent's rationality is bounded [10]. However, most coalition formation work does not consider any underlying structure between agents.

Recently there has been significant interest in social networks [12], [13] and the impact of those networks on performance of a group. For example, Xu looked at the impact of networks on coordination algorithms [14], [9], Kleinberg has looked at networks for search [15] and Boyd has looked at networks for *gossip*-based information dissemination [16].

VII. Conclusion and Future Work

Dynamically forming coalitions of sensors to acquire externally requested information in a complex environment is emerging as an important problem for large teams of sensors. To manage the computation and communication cost of forming these coalitions, networks are imposed over the sensors and coalitions are required to consist of sensors connected via that network. This paper presented new policies and results for adapting that network over time in order to allow the highest quality coalitions to form. One policy was shown to dramatically improve the team's ability to form high quality coalitions, while the other policy made it substantially worse. Moreover, an analysis of the network topologies resulting from both the new policies and from existing work showed that scale-free networks often formed but were not necessary for good performance.

Future work will look at addressing several short-comings of the current work. One key direction of future work will be to relate the networks between the agents to physical networks that exist within the team. These networks, while making the complexity of the adaptation algorithms higher, may improve performance overall since physical networks are likely to connect agents spatially close to one another and hence be more likely of collaborating on a task. Another direction is to consider cases where assets are not completely selfless, perhaps because they belong to different services or countries, and thus different coalition formation algorithms are required.

VIII. Acknowledgements

This research has been sponsored in part by AFOSR FA9550-07-1-0039 and AFOSR FA9620-01-1-0542.

References

[1] M. E. Gaston and M. desJardins, "Agent-organized networks for dynamic team formation," in *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*. New York, NY, USA: ACM, 2005, pp. 230–237.

[2] J. Yamamoto and K. Sycara, "A stable and efficient byuer coalition formation scheme for e-marketplaces," in *Proceedings of the 5th International Conference on Autonomous Agents*, 2001.

[3] C. Li and K. Sycara, "Algorithms for combinatirial coalition formation and payoff division in e-marketplace," in *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS 02)*, 2002.

[4] H. Moulin, "Cooperative microeconomics." Princeton University Press, 1995.

[5] O. Shehory and S. Kraus, "Formation of overlapping coalitions for for precedence ordered task-execution among autonomous agents," in *Proceedings of the 2nd International Conference on Multi Agent Systems (ICMAS)*, 1996.

[6] A. Mas-Colell, M. Whinston, and G. J., "Microeconomic theory." Oxford University Press, 1995.

[7] E. M. Arkin and R. Hassin, "On local search for weighted k-set packing," in *Proceedings of the 7th Annual European Symposium on Algorithms*, 1997.

[8] O. Shehory, K. P. Sycara, and S. Jha, "Multi-agent coordination through coalition formation," in *Agent Theories, Architectures, and Languages*. Springer Verlag, 1997, pp. 143–154.

[9] Y. Xu, P. Scerri, B. Yu, S. Okamoto, M. Lewis, and K. Sycara, "An integrated token-based algorithm for scalable coordination," in *AAMAS'05*, 2005.

[10] T. W. Sandholm and V. R. Lesser, "Coalition formation among bounded rational agents," in *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montréal, Québec, Canada, 1995, pp. 662–671.

[11] C. Li, S. Chawla, U. Rajan, and K. Sycara, "Mechanisms for coalition formation and cost sharing in an electronic market place," CMU, Tech. Rep., 2003.

[12] D. Watts and S. Strogatz, "Collective dynamics of small world networks," *Nature*, vol. 393, pp. 440–442, 1998.

[13] A.-L. Barabasi and E. Bonabeau, "Scale free networks," *Scientific American*, pp. 60–69, May 2003.

[14] Y. Xu, M. Lewis, K. Sycara, and P. Scerri, "Information sharing in very large teams," in *In AAMAS'04 Workshop on Challenges in Coordination of Large Scale MultiAgent Systems*, 2004.

[15] J. Kleinberg, "Complex networks and decentralized search algorithms," in *Proceedings of the International Congress of Mathematicians (ICM)*, 2006.

[16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE/ACM Trans. Netw.*, vol. 14, no. SI, pp. 2508–2530, 2006.