

Agent Organized Networks Redux

Robin Glinton and Katia Sycara and Paul Scerri

Robotics Institute
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213

Abstract

Individual robots or agents will often need to form coalitions to accomplish shared tasks, e.g., in sensor networks or markets. Furthermore, in most real systems it is infeasible for entities to interact with all peers. The presence of a social network can alleviate this problem by providing a neighborhood system within which entities interact with a reduced number of peers. Previous research has shown that the topology of the underlying social network has a dramatic effect on the quality of coalitions formed and consequently on system performance (Gaston & desJardins 2005a). It has also been shown that it is feasible to develop agents which dynamically alter connections to improve an organization's ability to form coalitions on the network. However those studies have not analysed the network topologies that result from connectivity adaptation strategies. In this paper the resulting network topologies were analysed and it was found that high performance and rapid convergence were attained because *scale free networks* were being formed. However it was observed that organizational performance is not impacted by limiting the number of links per agent to the total number of skills available within the population, implying that bandwidth was wasted by previous approaches. We used these observations to inform the design of a token based algorithm that attains higher performance using an order of magnitude less messages for both uniform and non-uniform distributions of skills.

Introduction

As robots and software agents become more capable, more robust and less expensive, multi-agent systems consisting of hundreds or thousands of agents are becoming possible (C., R., & B. 2005). Often, groups of agents within those systems will need to dynamically join together in a *coalition* to achieve a complex task that none of them can achieve independently. For sufficiently large, distributed organizations it can be infeasible for individuals to consider possible coalitions with any other possible agents, due to time, communication and computation constraints. One approach to overcoming this limitation is to impose some sort of network structure on the agents and require that agents only consider coalitions with neighbors in that network. In such a case, it

is clear that the nature of the network is critical to the quality of the coalitions that will be formed.

Gaston first proposed the network approach to overcoming the difficulties of coalition formation in very large organizations (Gaston & desJardins 2005a; 2005b). Critically, he showed how simple online adaptation of the network over time could dramatically improve the ability of the organization to form coalitions. However, Gaston did not report on the nature of the networks which were formed, limiting the ability to understand and improve on the results. Recent work in a variety of disciplines have shown how networks characteristics, such as a *small worlds* or *scale free* property (Watts & Strogatz 1998), can have a major impact on the performance of systems utilizing that network making an understanding of the topologies key to understanding how the algorithms work. The focus of this paper, is to understand the network structures evolved by Gaston's algorithm, the impact on the evolved structures of key environmental features and to use insights gained from that analysis to create more effective adaptation algorithms.

Analysis of the networks formed following Gaston's algorithm showed that *scale-free networks* were consistently being formed. In these networks, a small number of agents, referred to as *hubs*, were connected to a high percentage of the rest of the organization. These networks were intuitively effective because the hubs were very often able to form coalitions of themselves and a selection of the many agents they were directly connected to. However, in some real world systems, having agents connected to most of the rest of the organization is either infeasible or undesirable. To simulate such scenarios, we prevented the algorithm from exceeding a defined maximum number of links for any agent. We observed that as long as this maximum was at least equal to the required coalition size, there were no detrimental effects, but performance fell away quickly otherwise. Non-uniform distributions of skill throughout the organization were also problematic for scale free networks, because hubs were less often required as a member of a coalition.

Gaston's network adaptation algorithm relies on an agent knowing the current degree of all its neighbors' neighbors, which either heavily utilizes a point-to-point network or requires a broadcast communication medium. For many domains of interest, bandwidth will be tightly limited, thus it is of interest to determine whether it is possible to design

an algorithm utilizing less communication, but adapting the network as effectively. Moreover, analysis of the networks formed by Gaston’s algorithm showed that limiting the number of links an agent has, is generally a benefit to the team. A token-based algorithm was implemented to utilize point-to-point communication while exploiting the insights into the effective networks. The algorithm randomly walks a token from an agent wanting to change its links, to an agent to which a link would potentially be beneficial. The result was networks with a more even spread of links per agent. The performance of the resulting networks often out-performed Gaston’s algorithm, while using much less communication for adaptation.

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 $A = \{a_1, \dots, a_N\}$. 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 single skill given by $\sigma_i \in [1, \sigma]$ where σ is the number of available skills.

Agents must form coalitions on connected sub-graphs, to complete tasks introduced to the population. 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 skills R_{T_k} of size $|T_k|$. Skills are chosen uniformly from $[1, \sigma]$. Tasks are advertised for γ time steps. Any agents committed to a task are freed if the full compliment of skills 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 $a_i \in M_k$. The performance metric used throughout the remainder of this paper is:

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

where time is measured in discrete dimensionless units.

Coalition Formation

This section describes the algorithm used in (Gaston & desJardins 2005a) to form coalitions on a network. The pseudocode for this algorithm is given by Algorithm 1. The model of coalition formation is used in this paper is very simple, however, we are primarily interested in investigating alternatives to Gaston’s network adaptation strategies and need to use the same model of coalition formation to make a fair comparison.

On each iteration of the coalition formation algorithm each agent has the opportunity to adapt its local connectivity. Each agent chooses to adapt with probability $1/|A|$ where $|A|$ is the total number of agents.

Agents can be in one of three states, ACTIVE, COMMITTED, OR, UNCOMMITTED. An agent in the ACTIVE state is executing a task. An agent in the UNCOMMITTED state has not been assigned a task and an agent in the COMMITTED state has been assigned a task but the full com-

pliment of skills needed to accomplish the task are not yet available (complete coalition has not been formed).

If an agent elects not to adapt its local connectivity and it is in the UNCOMMITTED state, it has two options. The agent can initiate a coalition to fulfill a task that no other agent has committed to (lines 4-7 of Algorithm 1), or an agent can commit to a coalition if it is connected by a link to an agent already in that coalition and has a skill required by that coalition (lines 8-11 of Algorithm 1).

The probability that an agent will initiate a coalition is proportional to the number of its immediate neighbors in the UNCOMMITTED state. Formally:

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

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

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

```

JOINTTEAM()
(1) foreach  $T_k \in T$  in random order
(2)   if  $|M_k| = 0$  and  $s_i = UNCOMMITTED$ 
(3)      $r \leftarrow \text{UNIFORMRANDOM}([0, 1])$ 
(4)     if  $r < IP_i$ 
(5)       if  $\exists r \in R_{T_k} : r = \sigma_i$ 
(6)          $M_k \leftarrow M_k \cup \{a_i\}$ 
(7)          $s_i \leftarrow COMMITTED$ 
(8)     else if  $\exists a_j : e_{ij} = 1, a_j \in M_k$  and  $s_i = UNCOMMITTED$ 
(9)       if  $\exists r \in R_{T_k} : r = \sigma_i$  and r is unfilled
(10)         $M_k \leftarrow M_k \cup \{a_i\}$ 
(11)         $s_i \leftarrow COMMITTED$ 

```

Local Broadcast Algorithm

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 organizations ability to form coalitions. We briefly review that algorithm here, before analyzing the networks that result. Subsequently we refer to this as the local broadcast algorithm.

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:

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

The results reported in (Gaston & desJardins 2005a) 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.

Network Analysis

Our first objective in this investigation was to understand both the structure of the networks formed by the local broadcast algorithm and the relationship between the resulting network structure and the rate of task completion by coalitions formed on the network.

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.

To gain an understanding of the network structure formed by the local broadcast algorithm we conducted an experiment. Figure 1 shows a histogram of the links per agent, for the resulting network topology after the local broadcast 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 skills to complete a task with high probability. 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 skills uniformly distributed within the agent population. With a population of 500 agents this gives an expected number of 50 agents with each skill. 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 skills. This is important because it means that the hubs with high degree formed by the local broadcast algorithm waste bandwidth by communicating along many unnecessary links.

In the previous experiment we noticed that there was a large amount of variance in the results (not shown). The




Figure 1: The histogram of links per agent after 28000 iterations of the local broadcast algorithm, averaged over 50 trials.




Figure 2: Effect of limiting the max links/agent on performance of local broadcast.




Figure 3: Change in network adaptation performance of the local broadcast algorithm as d is varied.

hypothesis for this was that performance was sensitive to link density. Figure 3 shows the effect of the initial network link density on the performance of the local broadcast algorithm. The average link density increases with increasing d . Recall that initial networks are formed by randomly distributing agents within a unit square and connecting any two agents whose Euclidean separation is less than d . Notice, that even for a fixed value of d there is still significant variance, this is because the agents are distributed randomly, therefore the relationship between d and link density is not a direct proportion. Note that for large d performance is good because the initial networks formed already complete a large number of tasks per iteration. This means that there is only little room for improvement.

Figure 3 shows that as link density is increased (as d is increased) the performance of the local broadcast algorithm increases dramatically. Halving d results in a 4.5 times greater performance. However, a greater number of links results in higher bandwidth requirements for individual agents since they must communicate with all neighbors. In addition, this explains why the results of experiments in which performance is averaged over different values of d have a large amount of variance.

Token Algorithm

We used the insights gained through our study of the network topologies resulting from the local broadcast algorithm, to develop a more efficient token-based connectivity adaptation strategy. Specifically, we leverage two observations. The first is that the local broadcast algorithm was achieving good performance by forming networks with a relatively small number of agents that act as hubs and link most of the agents together. The second observation is that organization performance is not improved when the average number of links per agents is greater than the total number of available skills.

Consequently, we designed our algorithm to use a parameter MAX_DEGREE which limits the maximum number of links that an agent can have. Furthermore, we chose a token algorithm because token algorithms have been shown to collect as much information as broadcast algorithms while using much less bandwidth (Xu *et al.* 2005).

The pseudocode for the token algorithm is given by Algorithm 2. The token algorithm adapts local connectivity as follows: When an agent decides to form a new link it will break an existing link uniformly at random and instantiate a token. The token contains a running sum of the total number of links of each agent that has held the token thus far. The initiating agent passes the token to a neighbor using the $PASSBIASED(token)$ method given in line 3 of Algorithm 2. This method is described below. Upon receiving a token an agent a_j will generate a random real number $r \in [0, 1]$ and compare this to R_v where :

$$R_v = \frac{\text{number of links } a_j}{\sum_{a_i \in A_{visited}} \text{number of links } a_i}$$

where $A_{visited}$ is the set of agents visited by the token thus far. If $r < R_v$ the agent currently holding the token will destroy the token and form a new link to the agent that initiated the token. Otherwise the agent will pass the token on. The pseudocode for this approach is given by Algorithm 2.

Algorithm 2:

```

FINDTEAMMATETOKEN()
(1)  foreach  $t \in tokenList$ 
(2)    if  $A_{visited} = \emptyset$ 
(3)       $PASSBIASED(t)$ 
(4)    else if  $t.TTL \neq 0$ 
(5)       $t.TTL \leftarrow t.TTL - 1$ 
(6)       $r \leftarrow UNIFORMRANDOM([0, 1])$ 
(7)       $t.totalDegree \leftarrow t.totalDegree + this.degree$ 
(8)      if  $r < (R_v = \frac{degree}{t.totalDegree})$  and  $this.degree < MAX\_DEGREE$ 
(9)         $DROPEGE(t.creator, randomEdge)$ 
(10)        $CREATEEDGE(t.creator, this)$ 
(11)      else
(12)         $PASSBIASED(t)$ 

```

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.

Performance Analysis

We ran experiments to compare the performance of the local broadcast algorithm to the performance of the token algorithm under a variety of conditions. Figures 4,5,6 show a subset of the results of this investigation. Many others were omitted due to space constraints, however these results are representative.

Figure 4 was produced using a uniform distribution of skills within the agent population, a maximum of 5 links per agent and value of $d = 0.05$. The graph shows that

Algorithm 3: Algorithm used by agents to pass a token to a neighbor biased by agent degree.

```

PASSBIASED(token)
(1) totalDegree ← 0
(2) while true
(3)   n ← GETRANDOMNEIGHBORUNIFORM()
(4)   totalDegree ← totalDegree + n.degree
(5)   r ← UNIFORMRANDOM([0, 1])
(6)   if r <  $\frac{n.degree}{totalDegree}$ 
(7)     PASSTOKEN(token, n)
(8)   break

```

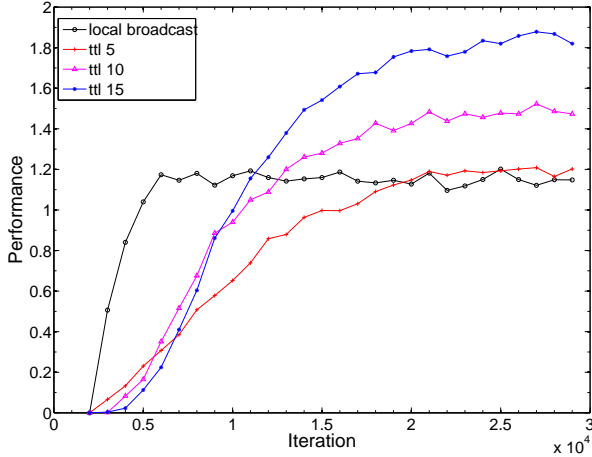


Figure 4: With a uniform distribution of skills and max of 5 links/agent local broadcast is outperformed for TTLs of 10,15, token algo uses an order of magnitude fewer messages.

the biased token algorithm outperforms the local broadcast algorithm for TTLs of 10 and 15. For a TTL of 15 the networks after adaptation complete an average of **60% more tasks**, with respect to the starting network configurations, than the local broadcast algorithm does. The biased token algorithm with TTL 15 and the local broadcast algorithm both use about 2×10^5 messages. This is to be expected because the maximum number of neighbors that a link can have is close to the expected number of neighbors considered by the token algorithm. The difference in performance can be explained by the relatively small maximum number of links allowed with respect to the total number of skills required to complete a task, 5 and 10 respectively.

Figure 5 was produced using a normal distribution of skills within the agent population, a maximum of 20 links per agent and with a value of $d = 0.15$. The graph shows that the biased token algorithm performs about the same as the local broadcast algorithm does. **However, the biased token algorithm uses 1.8×10^5 messages for a TTL of 15, while the local broadcast algorithm uses 1.6×10^6 messages an order of magnitude more.**

Figure 6 was produced by picking one skill to be 'rare'.




Figure 5: With a normal distribution of skills and a max of 20 links/agent, local broadcast is matched by the token algorithm for all TTLs using an order of magnitude fewer messages.

That is a single skill was distributed such that the expected number of agents in the population to have the rare skill was 2. A maximum of 20 links per agent was allowed and d was varied. The graph shows that the biased token algorithm outperforms the local broadcast algorithm for TTLs of 5, 10 and 15. For a TTL of 15 the networks after adaptation on average **double in performance** with respect to the local broadcast algorithm. Moreover, the token algorithm uses **an order of magnitude fewer messages. For a TTL of 15 the token algorithm uses 1.3×10^5 messages while the local broadcast algorithm uses 1.0×10^6 messages.** This dramatic increase in performance is due to the fact that the tokens allow an agent to search much deeper into the network for new agents to make connections with. This increases the probability of encountering agents with rare skills needed to complete tasks.

Figure 7 gives the histograms of links per agent of the final network topologies that correspond with the performance graphs of Figure 6. These graphs suggest that the more uniform distributions of links per agent, produced by the token algorithm, are responsible for the better performance. This makes sense in light of Figure 2, which showed that increasing the number of links beyond 10 (the number of skills) does not improve performance.

Related Work

Many authors have looked at coalition formation in multi-agent systems (Sandholm & Lesser 1995; Li *et al.* 2003; Shehory, Sycara, & Jha 1997). While some work, like the Gaston work built on here, focuses on protocols for efficient coalition formation (Gaston & desJardins 2005a), 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 rare

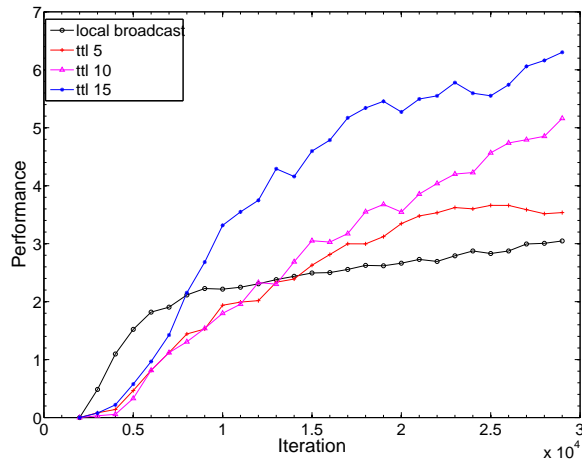


Figure 6: With some extremely rare skills and a max of 20 links/agent, variable d , local broadcast is outperformed for all TTLs of the token algorithm.

tionality is bounded (Sandholm & Lesser 1995). However, most coalition formation work does not consider any underlying structure between agents.

Recently there has been significant interest in social networks (Watts & Strogatz 1998; Barabasi & Bonabeau 2003) and the impact of those networks on performance of a group. For example, Xu looked at the impact of networks on coordination algorithms (Xu *et al.* 2004; 2005), Kleinberg has looked at networks for search (Kleinberg 2006) and Boyd has looked at networks for *gossip*-based information dissemination (Boyd *et al.* 2006).

Conclusion and Future Work

We determined that Gaston’s local broadcast network adaptation algorithm attains high performance and rapid convergence because it forms scale free networks, which are often impractical in real systems, because of the bandwidth burden on high degree hubs. However, we discovered that organizational performance is not impacted by thresholding the maximum number of links per agent to be at most equal to the number of skills available within the population. We used these observations to inform the design of a token based algorithm that attains higher performance using an order of magnitude less bandwidth than local broadcast. Furthermore, the token algorithm outperforms local broadcast for non-uniform distributions of skills, which is more typical in real systems. The network adaptation algorithms used here focused only on structure since Gaston found this to be the highly dominant feature. However, with non-uniform skill distributions, adaptations taking into account who has what skill may be useful.

The results presented here point to some interesting questions for the future. In this work a simple model of coalition formation was used, as well as a simple notion of coalition. More complex coalition types, e.g., super additive, or more

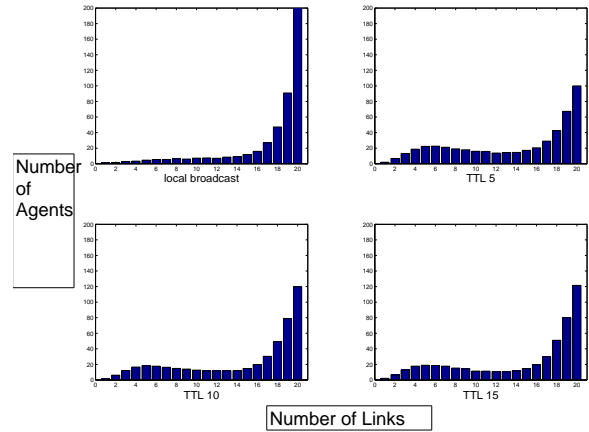


Figure 7: Histogram of network topology after 30000 iterations for the local broadcast algorithm and the biased token algorithm for TTLs of 5,10,15

sophisticated coalition formation protocols may prefer alternative network topologies. Our ongoing efforts are looking at these questions.

References

Barabasi, A.-L., and Bonabeau, E. 2003. Scale free networks. *Scientific American* 60–69.

Boyd, S.; Ghosh, A.; Prabhakar, B.; and Shah, D. 2006. Randomized gossip algorithms. *IEEE/ACM Trans. Netw.*

C., O.; R., V.; and B., M. 2005. Task inference and distributed task management in the centibots robotic system. In *AAMAS 05*.

Gaston, M. E., and desJardins, M. 2005a. Agent-organized networks for dynamic team formation. In *AAMAS ’05*, 230–237.

Gaston, M. E., and desJardins, M. 2005b. Agent-organized networks for multi-agent production and exchange. In *(AAAI)*.

Kleinberg, J. 2006. Complex networks and decentralized search algorithms. In *(ICM)*.

Li, C.; Chawla, S.; Rajan, U.; and Sycara, K. 2003. Mechanisms for coalition formation and cost sharing in an electronic market place. Technical report, CMU.

Sandholm, T. W., and Lesser, V. R. 1995. Coalition formation among bounded rational agents. In *(IJCAI-95)*, 662–671.

Shehory, O.; Sycara, K. P.; and Jha, S. 1997. Multi-agent coordination through coalition formation. In *Agent Theories, Architectures, and Languages*, 143–154.

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

Xu, Y.; Lewis, M.; Sycara, K.; and Scerri, P. 2004. Information sharing in very large teams. In *AAMAS’04*.

Xu, Y.; Scerri, P.; Yu, B.; Okamoto, S.; Lewis, M.; and Sycara, K. 2005. An integrated token-based algorithm for scalable coordination. In *AAMAS’05*.