

Comparing Market and Token-Based Coordination

Yang Xu
School of Info Sciences
University of Pittsburgh
Pittsburgh, PA 15260, USA
yxu@sis.pitt.edu

Paul Scerri and Katia Sycara
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
{pscerri, katia}@cs.cmu.edu

Michael Lewis
School of Info Sciences
University of Pittsburgh
Pittsburgh, PA 15260, USA
ml@sis.pitt.edu

ABSTRACT

Many coordination algorithms claim to be *general*, implying that they can be used to coordinate agents in a variety of domains. However, little work has been done to quantitatively compare distinctly different approaches to coordination across a range of domains. In this paper, we present a detailed comparison of two published coordination algorithms, performed in an abstract coordination simulation environment that allows extensive, quantitative experimentation. The simulator is used to compare two distinct approaches to coordination, token-based coordination and market based coordination. The results largely show the generality of different approaches, but performance and performance tradeoffs varies greatly across domains.

1. INTRODUCTION

Autonomous coordination is a complex process because several distributed algorithms are required to interact to produce agile, cohesive and efficient coordinated behavior. Because of the importance of autonomous coordination, many approaches have been developed, including approaches based on markets [1], tokens [5], etc. Typically, each of these approaches is designed to work in a specific domain, but the authors, usually with good reason, claim they will work in a wide range of domains. However, such claims are rarely quantitatively verified and, more importantly, competing approaches are rarely systematically compared. Thus, when a developer needs to select an approach to use in a particular domain, they are confronted with many claims but little concrete data with which to make a decision.

In this paper, we present an initial attempt at systematically and scientifically comparing distinct approaches to coordination. To make such a comparison feasible, we have developed an abstract simulation environment that is sufficient to capture a variety of real world concerns, but sufficiently abstract to be highly configurable and very fast. We used the coordination simulator to investigate the relative strengths of three distinct approaches to coordination:

auction-based coordination [2]; token-based coordination [5]; and a hybrid of the two.

Based on an analysis of previous literature [1, 5], several hypotheses can be formed about the relative performance of the algorithms. Auctions are focused on maximizing overall utility taking into account the *bids* of all team members [1]. Token-algorithms are focused on scalability, hence they minimize communication, sometimes at the expense of overall utility. Thus, the clearest hypothesis is that auctions will communicate more than token algorithms, but result in better allocations of tasks and resources. More subtly, the performance advantage of an auction should be most pronounced when small changes in allocations lead to big differences in performance, i.e., typically highly constrained cases, while the token algorithms should maximize their communication advantage when the probabilistic models they rely on are most advantageous, i.e., weakly constrained cases. The empirical results support these hypotheses. Initial experiments suggested that auctions find superior allocations because they compare many options, while tokens use little communication by quickly focusing on the agents most likely able to perform tasks or having most use for resources. A hybrid algorithm that uses tokens to solicit auction bids from those agents most likely to submit winning bids then uses an auction to select from between the small number of bids should perform well. However, this hybrid algorithm should only perform well under restricted circumstances. We implemented the hybrid algorithm and compared its performance to the other algorithms.

2. COORDINATION ALGORITHMS

In this paper, we only consider the coordination tasks that can be performed by both tokens and auctions and is limited to task and resource allocation. The coordination problem in this paper is defined as: Agents, $A = \{a_1, \dots, a_k\}$, are cooperating on a joint goal G . G is broken into discrete subtasks $\alpha_1, \dots, \alpha_n$. Each subtask α_i when is applicable, must be performed by a single agent a to receive reward. The amount of rewards is defined as $Reward(Cap(a, \alpha), Holds(a)) \rightarrow \mathcal{R}$ where it depends on the capacity ($Cap(a, \alpha)$) that agent a can perform α and the exclusive resources that are available for a to perform α , i.e., a bucket of water is required for distinguishing a fire. The function $Assigned(a, \alpha) = 1$ if agent a is assigned to task α , otherwise it is equal to 0. Only one agent may be assigned a task at any time, i.e., $\sum_{a \in A} Assigned(a, \alpha) \leq 1$. The coordination is to maximize the reward to the team, while minimizing the *costs of coordination*. The overall reward is:

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'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

$$\sum_{i=0}^n \sum_{a \in A} \text{Assigned}(a, \alpha_i) \text{Reward}(\text{cap}(a, \alpha_i), \text{Holds}(a))$$

We are specifically concerned with only the volume of Messages as costs of coordination. It is the number of times agents communicated, either between themselves or with the auctioneer.

The first of the algorithms we used for comparison is market-based approach. Our implementation of this approach was based on TraderBots [2] with adaptations where are necessary to make a comparison possible. In our market-based approach, one agent acts as auctioneer and both tasks and resources are treated as merchandise. Agents bid for either single items or combinatorial sets of items in order to maximize their own utilities. The auctioneer maximizes its utility by "selling" their "merchandise". In this approach, Sandholm's winner determination algorithm [3] is used to determine the allocation for tasks and resources by the auctioneer. Because of the centralized position of the auctioneer, it develops a complete knowledge of how agents will use a task or resource if allocated. Thus, the auctioneer can perform assignments that maximize the team utility. To be fair to all the bidders, the auction should last for a fixed period of time. Where early determination is infeasible; Agents are allowed to bid for resources after tasks have been allocated. Moreover, to prevent deadlock in resource allocation, agents are only allowed to bid for resources for their *first* pending task.

In Token-based algorithms [5], tokens encapsulate both resources and tasks, allowing actors to locally decide what to do with the token. For example, a *task token* allows an actor to decide whether to perform the task or pass it off for another actor. An intelligent routing algorithm is built in the token-based approach to help agents build local decision theoretic models to determine when and where to pass tokens. By utilizing the relevance between tokens, intelligent routing algorithm is able to efficiently deploy tokens to make higher utility with less communication [5]. In this paper, tasks are allocated by the LA-DCOP token algorithm [4], agent is allow to reject previous accepted task but accept another task that it can get more reward.

The hybrid algorithm works in the following way. The auctioneer algorithm runs exactly as before, except that instead of broadcasting announcements for auctions an *auction token* is created. Each auction token is allowed to exist from the starting of the auction to the end of the auction being closed. The auctioneer has a probabilistic model of the team state, just as all agents do in the token-based approach. The auction token is then intelligently routed to the agents most likely to be able to submit the best bids. The token stops moving after the auction it presents is closing or has visited a fixed number of teammates. Note that although the intelligent routing algorithm should work to route tokens for higher bids, it should not work better than the token-based approach because of incapability of making use of the relevance between tasks and resources which have been encapsulated into auction tokens. The auctioneer determines the winner of the auction and allocates tasks and resources the same as in the basic auction case. We expect that the hybrid approach should keep high rewards as auction-based approach but reduce communication over the basic auction, by targeting only those agents likely to make




Figure 1: The reward for heterogeneous teams dramatically decrease in auction and hybrid approaches but maintains constant in token-based approach where more messages are required.

good bids and reduce computation by limiting the number of bids the auctioneer must deal with.

3. EXPERIMENTS

We implemented the three approaches in an abstract simulator called CoordSim. If not otherwise stated, the CoordSim are configured as follows. There are 100 agents to perform 50 tasks with 50 resources. Each task requires only one resources which could be interchangeable with four others. In the default setup, there is only one type of capability required and all agents have none-zero value for this capability, i.e., all agents are at least somewhat capable of all tasks. Auctions are held open for 40 time steps and the task tokens, resource tokens are allowed to move unless accepted. The initial threshold on a task token is 100, meaning that the task will not be accepted by an agents until it can get a reward more than 100 by performing this task. Simulation runs for 2000 steps and the results are based on 100 runs.

3.1 Heterogeneous Team

In the first experiment, we examined team performance by varying team composition and the capabilities required to perform tasks. For example, in an emergency response experiment some agents might only be able to fight fires while others could only provide medical treatment. In this experiment, we varied the number of capabilities from 1 to 46 where fewer agents in the team were available to perform particular tasks and in the most heterogeneous condition, only two agents on average are capable of performing a task.

The experimental results in Figure 1 show that for heterogeneous teams, auction and hybrid approaches earn less reward as the team becomes more heterogeneous because there are fewer agents able to compete for the more specialized tasks. The advantages of teamwide maximization of utility by the auctioneer decrease as there are progressively fewer feasible alternative bids. In contrast, reward

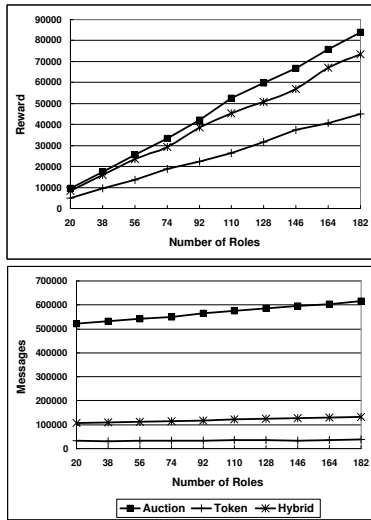


Figure 2: Reward and Messages increase dramatically with the number of tasks.

for the token-based approach remain almost flat with increasing specialization. We propose two reasons. One is that token-based approach greedily finds reasonable solution rather than searching for the optimal. As the other reason, by passing a higher number of tokens around the network and making use of the relevance between them, intelligent routing algorithm gets better knowledge to route tokens. This is manifested that although the average distance to route a token increases with heterogeneity as reflected in an increase in message, token-based approach maintains the same level of reward.

3.2 Time Critical Tasks

In the second experiment, we investigated team performance when many tasks needed to be performed within a short period of time. To increase their reward, teams were required to perform tasks and allocate resources as rapidly as possible. In this study we varied the number of tasks the teams were required to finish from 20 to 182. After 2000 time steps, the accumulated reward and message count were recorded as shown in Figure 2.

All three approaches performed more tasks in order to get higher reward. As expected, the auction approach attained higher reward than the hybrid or token-based approaches. Considering both reward and messages, however, the hybrid approach performs well by almost matching the reward obtained by the auction at just a quarter of the communication cost. The reason why the hybrid approach achieves such good performance with so little communication overhead is that the intelligent routing algorithm limits communication to a small number of agents while high bidders must always be informed in auctions.

3.3 Competitive Resources

The third experiment used 200 tasks each requiring an average of four resources with no interchange possibilities. As available resources are increased from 4 to 40, competition for them declines and they become less likely to be a bottleneck.

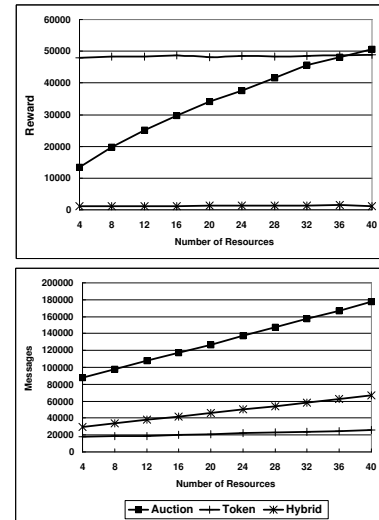


Figure 3: Reward decreases with scarcity in auction approach and is very low in hybrid approach, but with token-based approach is uniformly high

The experiment was stopped after 1000 time steps. Figure 3 shows that the reward for the auction based approach increased rapidly with increases in resources. Both the token-based and hybrid approaches remained flat with token-based approach earning the highest reward at all levels of scarcity while the hybrid approach yielded very limited reward. We hypothesized that because resource contention in this experiment was high the centralized control of the auction and hybrid approaches would often force agents to either bid for all four resources together or miss the task while the distributed token-based approach weakened this constraint.

Acknowledgements

This research has been sponsored in part by AFRL/MNK Grant F08630-03-1-0005 and AFOSR Grant F49620-01-1-0542.

4. REFERENCES

- [1] B. Dias, R. Zlot, N. Kalra and A. Stentz. Market-based multirobot coordination: A survey and analysis. In *Technical report, CMU-RI-TR-05-13*, 2005.
- [2] B. Dias and A. Stentz. Traderbots: A market-based approach for resource, role, and task allocation in multirobot coordination. In *Technical report, CMU-RI-TR-03-19*, 2003.
- [3] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. In *Artificial Intelligence*, 135, 2002.
- [4] P. Scerri, A. Farinelli, S. Okamoto and M. Tambe. Allocating tasks in extreme teams. In *Fourth Int. Conf. on Autonomous Agents and Multiagent Systems*, 2005.
- [5] Y. Xu, P. Scerri, B. Yu, S. Okamoto, K. Sycara and M. Lewis. An integrated token-based algorithm for scalable coordination. In *Fourth Int. Conf. on Autonomous Agents and Multiagent Systems*, 2005.