

An empirical evaluation of auction-based task allocation in multi-robot teams

(Extended Abstract)

Eric Schneider¹, Ofear Balas², A. Tuna Özgelen³, Elizabeth I. Sklar¹ & Simon Parsons¹

¹Dept of Computer Science
University of Liverpool
Liverpool, UK

²Dept of Computer & Information Science
Brooklyn College
City University of New York, USA

³Dept of Computer Science
The Graduate Center
City University of New York, USA

{eric.schneider,e.i.sklar,s.d.parsons}@liverpool.ac.uk,
ofear.balas@gmail.com,aozgelen@gc.cuny.edu

ABSTRACT

Task allocation is an important topic in multiagent and multi-robot teams. In recent years, there has been much research on the use of auction-based methods to provide a distributed approach to task allocation. Team members bid on tasks based on local information, and the allocation is based on these bids. The focus of prior work has been on optimal allocations and has established that auction-based methods perform well in comparison with optimal methods, with the advantage of scaling better. Here we take a different approach, comparing auction-based methods not on the optimality of the allocation, but on the *efficient execution* of the allocated tasks. This approach factors in aspects such as the utilisation of the team members and the degree to which they interfere with each others' progress, giving a fuller picture of the practical use of auction-based methods.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Coherence and coordination, multiagent systems*

General Terms

Experimentation, performance

Keywords

Multi-robot team, auction mechanisms, task allocation

1. INTRODUCTION

This paper is concerned with the *multi-robot routing* problem. This is a task allocation problem in which multiple robots are required to reach particular *target* locations, and each target must be visited by one robot only. The problem is computationally hard: the number of ways of allocating m targets to n robots quickly defeats attempts to use standard optimisation techniques. This has led researchers to look for more efficient solutions [1, 3, 4]. For example, [3] applied *auction mechanisms* to allocate the target points and

demonstrated that multi-round auctions performed experimentally close to optimal allocations. Later work considered the *sequential single-item* auction [2] as an alternative to combinatorial auctions for multi-robot routing.

Our extension to this work evaluates the effectiveness of auction mechanisms using metrics that measure how tasks are executed in practice. This reveals the trade-offs between mechanisms in different scenarios, and allows us to draw conclusions about the applicability of those mechanisms.

2. APPROACH

Previously, we developed the *HRTeam* environment in which to conduct experiments with multi-robot teams, both in simulation and with physical robots [5]. The HRTeam system logs data that we use to compute a range of performance metrics, such as distance travelled by individual robots and overall runtime to complete a multi-robot routing mission. We measure the *distance travelled* as the actual distance moved by the robots rather than the shortest distances between the target points. As Figure 1 shows, the resulting paths show plenty of variation from straight line paths. We also measure *run time*, the time between the start of an experiment and the point at which the last robot on the team completes the tasks allocated to it. This includes both the *deliberation time*, the time to determine which tasks are allocated to which robots, as well as the execution time, the time during which robots travel to points.

We considered four different mechanisms for task allocation. In *round robin* (RR), the first target point is allocated to the first robot, the second point to the second robot and so on. In an *ordered single item auction* (OSI), each robot makes a bid for the first point, where the bid is the distance that the robot estimates (using an A* path planner) it will have to travel from its current location. The point is allocated to the robot that makes the lowest bid and the remaining points are auctioned in the same manner. In a *sequential single item auction* (SSI) [2], all of the target points are presented to all the robots simultaneously. Each bids on the target point with the lowest cost and the point with the lowest bid is assigned to the robot that made the bid. The remaining points are then re-auctioned, until all points have been allocated. A *parallel single item auction* (PSI) [2] starts like SSI with all robots bidding on all points, but all of the target points are allocated in one round, with each point going to whichever robot made the lowest bid on it.

Appears in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, Lomuscio, Scerri, Bazzan, Huhns (eds.), May, 5–9, 2014, Paris, France.

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: Example paths taken by robots

3. RESULTS

We conducted experiments using two different starting configurations for 3 robots and four different task point configurations, comparing each of the four different allocation mechanisms mentioned above. Each was run multiple times. Representative results are presented here. Note that robots visited task points in the order in which they were allocated, which wasn't necessarily the shortest path amongst the complete set of points allocated to a given robot. Figure 1 shows paths taken by the robots in runs across different scenarios as solved by different allocation mechanisms. The propensity for SSI to produce tight groupings is clear in its handling of the distributed start points. Figure 2 plots the average values of total distance traveled, deliberation and run times over 10 runs for each of the four mechanisms, covering both the clustered and distributed start points.

Overall, our analysis supports the results in [2], showing the effectiveness of SSI in finding solutions to the multi-robot routing problem when the overall distance covered is the most important performance metric. For both the clustered and distributed start points, SSI generated solutions which required the team to travel the smallest combined distance on average, and thus produced the shortest run time. As [2] points out, PSI can come up with arbitrarily poor allocations because it does not take synergies between target points into account. It can also skew the distribution of tasks between robots. Figure 1 (upper right plot) reveals that in the clustered start case, it allocates all the target points to one robot. This skew means that although PSI is not much worse than SSI or OSI on overall distance, the runtime for PSI is more than twice that for SSI and OSI.

The one area in which SSI performs worse than other mechanisms, particularly PSI, is in deliberation time. For the scenarios we consider here, the cost of carrying out the allocation is negligible, with the deliberation time being less than 1% of the total time for completing the mission. The total number of bids to allocate m tasks to n robots is: $n * (m(m+1))/2$ and it is conceivable that this could become problematic. For example, consider allocating 500 targets to 100 robots: SSI would require over 12 million bids; enough to make deliberation time a significant contributor to the time for task completion. This might make OSI or PSI worth considering for larger deployments.

4. SUMMARY

This paper has studied the performance of a number of auction mechanisms on a version of the multi-robot routing

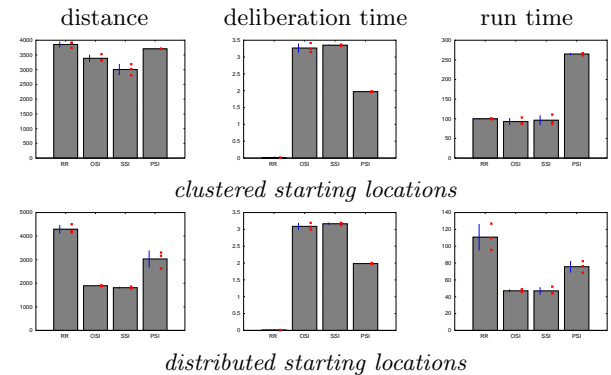


Figure 2: Metrics: mean, standard deviation (blue), distribution (red); column order: RR, OSI, SSI, PSI.

problem, measuring how the mechanisms will perform on a fleet of real robots. The main result is that the sequential single item (SSI) auction broadly outperforms other single item auctions by our metrics, though it does not perform best on all of them for all scenarios and the order in which allocated points are visited is a significant factor (discussion of this aspect is beyond the scope of this short paper). However, there do seem to be trade-offs, especially in terms of the total number of bids required by the mechanisms. This suggests that the SSI auction might have issues with scaling to larger routing problems than we study here, especially if communication bandwidth is restricted. In other words, the high performance of the SSI comes at a cost that might be hard to pay for some scenarios. Other mechanisms we tested, which can scale better, might be preferable on such scenarios despite their poorer performance. This topic is something we plan to study next in more detail.

Acknowledgements

This work was partially funded by the National Science Foundation grants #IIS-1116843 and #IIS-1338884, a University of Liverpool Research Fellowship and a Fulbright-King's College London Scholar Award.

5. REFERENCES

- [1] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: A survey and analysis. Carnegie Mellon Univ Tech Report CMU-RI-TR-05-13, 2005.
- [2] S. Koenig, C. Tovey, M. Lagoudakis, D. Kempe, P. Keskinocak, A. Kleywegt, A. Meyerson, and S. Jain. The Power of Sequential Single-Item Auctions for Agent Coordination. In *Proc of AAAI*, 2006.
- [3] M. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. Kleywegt. Simple auctions with performance guarantees for multi-robot task allocation. In *Proc of IROS*, 2004.
- [4] M. Mataric, G. Sukhatme, and E. Ostergaard. Multi-robot task allocation in uncertain environments. *Autonomous Robots*, 2003.
- [5] A. T. Özgele, E. Schneider, E. I. Sklar, M. Costantino, S. L. Epstein, and S. Parsons. A first step toward testing multiagent coordination mechanisms on multi-robot teams. In *ARMS Workshop*, 2013.