A first step toward testing multiagent coordination mechanisms on multi-robot teams

A. Tuna Özgelen[1], Eric Schneider[2], Elizabeth Sklar[1,2], Michael Costantino[3], Susan L. Epstein[4], and Simon Parsons[1,2]

[1]The Graduate Center, [2]Brooklyn College, [3]College of Staten Island, [4]Hunter College
The City University of New York, New York, USA
tunaozgelen@gmail.com, nitsuga@pobox.com, sklar@sci.brooklyn.cuny.edu, michael.costantino@cix.csi.cuny.edu, susan.l.epstein@gmail.com
parsons@sci.brooklyn.cuny.edu

Abstract. Our work investigates the effectiveness of multiagent coordination mechanisms, not for groups of software agents, but rather for teams of physical robots. The long-term aims are to identify which multiagent coordination mechanisms are robust enough to work on physical robot teams, despite the inherent problems with noise and error, and to identify which of those mechanisms work best for particular types of missions. This paper takes a first step in this direction, by comparing the performance of a simple auction mechanism for task allocation to a greedy mechanism, as measured on a physical multi-robot team.

1 Introduction

The field of multiagent systems (MAS) has produced much work on coordination mechanisms, methods that could potentially be applied to multi-robot systems (MRS). Few of these models, however, have been evaluated on a team of real robots. Most are either evaluated theoretically, or are evaluated only in simulation. While theoretical and/or simulated evaluation is necessary to the development of any such model, neither tells us whether the model will effectively support a physical robot team, with all the practical issues that real robots entail. In short, while the literature reports that some coordination mechanisms are better than others in non-robot settings, we do not know whether they are robust enough to be effective on robots that operate in the physical, rather than a simulated, world. Our long-term goal is to address this gap in the literature, and we aim to evaluate a range of multiagent coordination mechanisms on physical multi-robot teams. The results presented here represent a first step toward this goal and demonstrate the effectiveness of a simple auction over a greedy mechanism for task allocation as measured on a physical multi-robot team.

Our work has two characteristic features. First, we take what we call a "rough and ready" approach. We use a team of small, inexpensive, off-the-shelf robots with limited functionality, and work within the limitations that such equipment imposes, rather than engineer the problems away by using more powerful robots.

We believe this makes our results more applicable—if we find an interesting signal through the noise associated with our low-end robots, others will find a similar signal through the lesser noise associated with higher-end robots. It also makes our work closer to deployment. Although our physical environment is still far from the real environment our team will eventually inhabit, it is closer than a precision-built, highly-instrumented laboratory. The second characteristic feature of our work is the types of missions that our multi-robot team addresses. The missions we study are abstract versions of those from *urban search and rescue* (USAR) [12, 21] and *humanitarian de-mining* [11, 23]. Both these applications require a team of robots to move through a space and search it, either looking for victims (in the case of USAR), or mines (in humanitarian demining). A simple version of this type of mission is our current focus, and the one we consider here.

The remainder of this paper is structured as follows. Section 2 provides some background on coordination mechanisms, particularly market-based approaches, since one of the mechanisms we test here is an auction. Section 3 motivates the work. Section 4 describes the experiments we carried out, with results and some discussion. Conclusions appear in Section 5.

2 Background and Related Work

Prior research on market-based approaches to multi-robot coordination draws on Smith's Contract NET Protocol [28] designed for for software agents and Wellman and Wurman's [30] concept of *market-aware agents*. Wellman and Wurman [30] state that every decision in a MAS is about resource allocation. Thus, market-based approaches have been widely used in multi-robot systems to optimize resource usage, communication methodology and task-completion time [9]. This section summarizes some of the previous work in this area.

A primary strength of market-based approaches is their reliance on local information or the self-interest of the agents to arrive at efficient solutions to large-scale, complex problems that are otherwise intractable [2, 3]. The most common instantiations of market-based approaches in MAS and MRS are auctions. Auctions are commonly used for distribution tasks, where resources or roles are treated as commodities and auctioned to agents. Based on their private preference information for a particular commodity, agents bid in these auctions. A significant body of work analyzes the effects of different auction mechanisms [1, 15, 16, 28, 31], bidding strategies [18, 24, 29], dynamic task re-allocation or swapping [10] and levels of commitment to the contracts [19] on the overall solution quality. The application domains vary from loosely-coupled tasks, such as exploration [9, 17, 31], to tightly-coupled tasks [9, 13] such as box pushing and formation control [20], which require close coordination between robots. Auctions have also been used in such contexts as role allocation in robot soccer [6] and multi-robot routing [18].

In domains where there is a strong synergy between items, single-item auctions can result in suboptimal allocations [1]. In the multi-robot exploration domain, the domain of interest here, strong synergy may exist between target

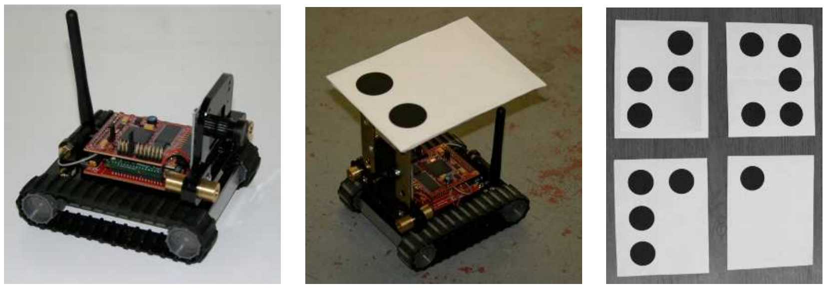

(a) Unmodified (b) With hat (c) Braille hats

Fig. 1. The Surveyor SRV-1 Blackfin

points that robots need to explore. Combinatorial auctions can remedy this limitation by allowing agents to bid on bundles of items and by attempting to minimize the total travel distance, because they can take into account the relative proximities of target locations [14]. However, combinatorial auctions suffer from the computational costs of exponential bid generation and bundle valuation by the agents, and winner determination by the auction mechanism itself, all of which are NP-hard [15]. Sequential single-item auctions (SSA), which are multi-round auctions, have been studied as an alternative to combinatorial auctions [14, 15]. In SSA, at every round the agents rebid on all unallocated tasks and only the auction which received the lowest bid is cleared.

3 Motivation

Our long-term aim is to test a range of coordination techniques from the MAS literature, to see how techniques tested theoretically and in simulation perform in the rough-and-ready world of low-end physical robotics. This paper concentrates on the use of a simple market-based mechanism applied to a specific task: efficient exploration of a region with a team of autonomous robots. It is natural to operate several robots in parallel, but this raises the question of how best to coordinate them. We studied a simple version of this problem. Given a set of n robots, we considered how best to have them visit m different positions, which we call *interest points*. The specification of interest points is an abstraction for the allocation of search areas to robots; we assume that by visiting a point of interest, a robot has searched the relevant area. Future work will additionally evaluate the follow-on task of actually searching the area, but currently we focus on just getting there.

The baseline allocation method we consider here is a greedy mechanism that takes each robot in turn and assigns it an interest point, in the same way that taxi ranks are typically handled at airports. Both taxis and customers (riders) form a queue, and the first customer is assigned to the first taxi, the second customer to the second taxi, and so on. If the number of taxis exceeds the number of customers, then the surplus taxis wait in the queue until more customers arrive; similarly if the number of customers exceeds the number of taxis, the surplus customers wait. It might be satisfactory in a multi-robot system where the robots all start in roughly the same location and their destinations are evenly distributed (as for airport taxis in the average case), it has clear pitfalls in cases where the distributions of numbers and locations of robots and/or interest points are lumpy. For example, if $m > n$, where some robots would have to explore several points, a balanced mechanism should allocate points so that all robots travel approximately the same total distance.

Here we test the greedy-taxi mechanism against a simple auction mechanism that attempts to balance, across the team, the robots' estimated costs to complete all tasks. Given robots' initial locations and a list of interest-point locations, robots bid for interest points. Bids are determined by the robots' distance to the points, as calculated by an A* path-planner with a map of the area. Points are allocated one at a time, in a sequence of auctions. In each auction, a robot considers its last position in its bid for its next point; for its first point, this will just be its starting location, whereas for subsequent points, this will be the last point it "won" in a previous auction. Such a robot estimates the cost to travel from its most-recently-allocated interest point to the new point. A robot not yet allocated any points estimates its travel cost from its initial position, before the start of the first auction. In some situations, this approach will clearly be less efficient in total distance traveled than a combinatorial auction which allowed robots to bid on bundles of locations. The simple auction, however, will likely be more efficient in terms of computational effort, given the well-documented computational cost of combinatorial auctions [7, 22], and hence more practical in a real-time, dynamic environment. This paper tests the hypothesis that the simple auction mechanism provides better task allocations than the greedy-taxi approach.

4 Experiments

In this section, we describe our experimental environment and how we measured the relative performance of the simple auction and the greedy-taxi mechanisms for task allocation.

4.1 Experimental Environments

Our experimental environment has been documented before [26, 27], and so is described only briefly here. Our multi-robot team is comprised of inexpensive, limited-function platforms. For these experiments, we used the Surveyor SRV-1 Blackfin[1], a small tracked platform equipped with a webcam and 802.11 wireless. Because the Blackfin (shown in Figure 1(a)) has very limited on-board

[1] `http://www.surveyor.com/SRV_info.html`

Fig. 2. The experimental arena

processing, our robots rely on an off-board processor[2] for intelligent robot control. This *robot controller* instantiates a Player [8] client to interact with the physical robot hardware. The robot controller formulates bids for the simple auction and performs path-planning on the set of points allocated to the robot, to determine where the robot should go next. Movement is facilitated using a cognitive architecture called *SemaFORR* [5], a FORR-based decision system [4].

As described in [27], this setup is complemented by a simulator[3] using Stage [8], which allows us to run experiments in parallel physical and simulation environments. The same software controls the robots in both cases; the robots (real or virtual) simply operate in different environments.

The robot team is deployed in an arena with a floor plan that imitates the layout of an office building, with six rooms that open onto a central corridor. Figure 2 is a picture of the arena. Localization for the robot team is provided by a set of six *camera controllers*. Each camera controller talks to one webcam that is suspended above the robots' physical arena and identifies the robots in its field of view. To distinguish amongst the robots, each robot is equipped with a unique hat (e.g., Figure 1(b)) depicting one letter from the Braille alphabet (Figure 1(c)).

[2] The physical results presented here were run on three octocore Xeon 2.66GHz processors, operating Ubuntu 12.04, and executing the following processes: a central server, an auction manager, 4 robot controllers (Player clients), 4 robot drivers (Player server), and 6 camera controllers. Communication between processors was conducted over standard ethernet; communication with robots was conducted using 802.11 wireless on a closed local area network.

[3] The simulation results presented here were run on one quadcore Intel 3.3GHz processor, operating Ubuntu 12.04, and executing the following processes: a central server, an auction manager, 4 robot controllers (Player clients), 4 robot drivers (Player server), and the Stage simulator.

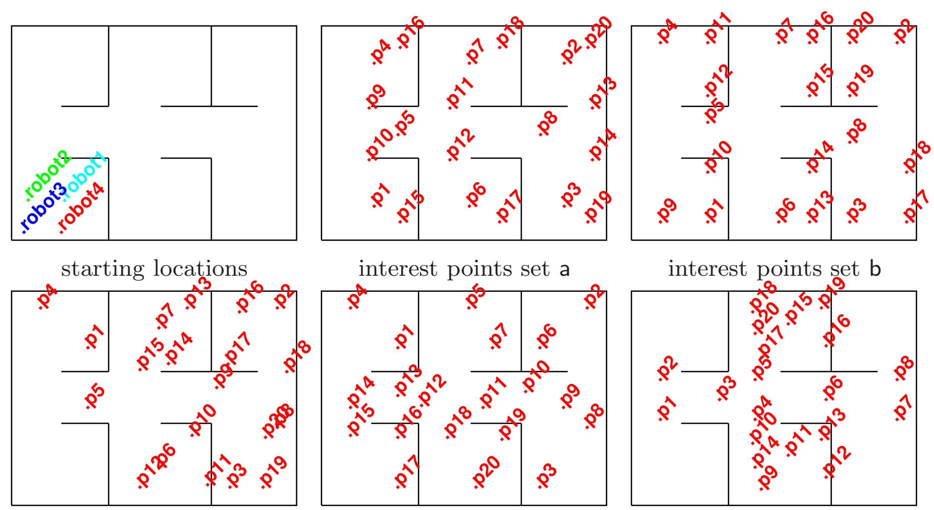

starting locations | interest points set a | interest points set b

interest points set c | interest points set d | interest points set e

Fig. 3. Scenario definitions

4.2 Experimental Setup

As described earlier, the team's mission is to explore the space, by visiting a number of interest points. A *scenario* is a mission defined by a specific set of parameters: the number of robots on the team (n), the starting locations for the robots, the number of interest points to visit (m), and the locations of the interest points. Thus, a scenario can be described by a tuple

$$\langle n, \{(x_0, y_0), \ldots, (x_{n-1}, y_{n-1})\}, m, \{(x'_0, y'_0), \ldots, (x'_{m-1}, y'_{m-1})\}\rangle$$

The experiments reported here measured results in five different scenarios. All scenarios involved $n = 4$ robots, $m = 20$ interest points, one fixed set of starting locations (Figure 3, top left), and five different sets of interest point locations (labelled a through e, Figure 3). Experiments were conducted with each scenario using the two task allocation mechanisms discussed above (simple auction and greedy-taxi, labelled A and G, respectively), five times in the physical environment and five times in the simulated environment (labelled P and S, respectively). Thus 100 experimental trials were conducted in all:

$$(5 \textit{ scenarios} \times 2 \textit{ allocation mechanisms} \times 2 \textit{ environments}) \times 5 \textit{ trials}$$

Each experiment recorded:

Run time: the elapsed time between the start of an experiment and the time that the last robot completed the tasks allocated to it. Run time varies with the scenario and the way that tasks are distributed amongst robots.

Deliberation time: the elapsed time between the start of an experiment and the completion of the task allocation mechanism, i.e., the point at which the robots began to execute tasks (move to the interest points allocated to them). Deliberation time varies with the scenario and the allocation mechanism.

Near collisions: with several robots moving around a limited space, especially one that requires robots to move through a single corridor to navigate from room to room, robots naturally get in each other's way and are in danger of collisions. The current version of our system detects situations where a collision is likely, stops one robot, and gives the other the right of way. Since such episodes impact the other metrics, we count these episodes.

In addition, for each robot we recorded:

Idle time: some robots complete their tasks before other robots do. Idle time is the time that elapses between when a robot completes its last task and the time that every robot in the team completes its last task. Idle time varies from one robot to the next; it is affected by the way that tasks are distributed amongst the robots and also by the scenario. It provides an estimate of the spread of tasks between robots.

Delay time: the total time the robot spends stationary as a result of collision avoidance. Delay time varies from one robot to the next, as well as with the way the tasks are distributed.

Distance: the total length of the path traveled by the robot, i.e., the sum of the differences of consecutive position updates (Euclidean distance). Note that this is not completely accurate for the physical robots, because their positions are determined by the overhead cameras, and any errors in a robot's position (e.g., due to mis-identification of a robot hat) are included in the distance calculation.

From these values we also compute the **speed** of each robot:

$$speed = \frac{distance}{travel\ time - delay\ time}$$

where travel time minus delay time is the amount of time that the robot is actually moving.

4.3 Results

This section describes our results and draws some conclusions from them. Section 4.4 provides more detailed discussion.

Figure 4 illustrates the allocation of interest points to robots for all 50 experiments conducted in the physical environment. Each row in the chart represents one interest point. Each column in the chart represents one trial. The cells (row-column intersections) are color-coded for each of the 4 robots used in the experiments. The experiments are displayed in order, starting with 5 trials using the simple auction allocation mechanism in scenario `a`, and ending with 5 trials using the greedy-taxi allocation mechanism in scenario e. The chart shows that for the simple auction mechanism trials, in the left half of the graphic, the allocation was fairly consistent from one trial to another; whereas the greedy-taxi mechanism was more random (in fact, a round-robin ordering was used to rotate which robots were offered points in which order).

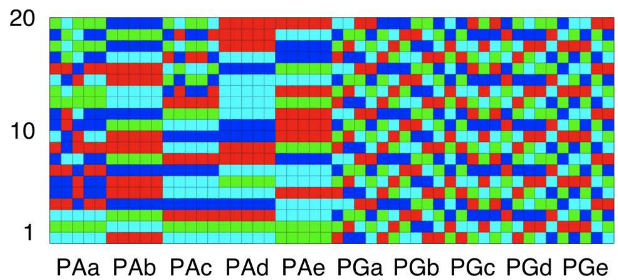

Fig. 4. Distribution of interest points for all experiments in physical environment. See text for explanation. Legend: robot1=cyan, robot2=green, robot3=blue, robot4=red.

Figure 5, which considers the metrics summed over all the different experimental scenarios, provides the headline results. Since our main interest is in the performance of coordination mechanisms on physical robots, what we are most interested in is the comparison between the PA (physical robots, task allocation by simple auction) and PG (physical robots, task allocation by greedy-taxi) results. These show that although the auction takes slightly longer to provide an allocation (deliberation time, Figure 5(b)), on all other metrics auction allocation outperforms the greedy-taxi allocation mechanism. Since the auction considers the distance that the robots travel, one might well expect that the runtime (Figure 5(a)) for PA would be better than the runtime for PG. More interestingly, the idle time (Figure 5(c)) also improved, which suggests that the tasks are more evenly distributed. (Recall that idle time shown here is the sum across the entire team.) In addition, the auction leads to less delay time (Figure 5(d)) and fewer near collisions (Figure 5(e)). (Delay time and near collisions are related since a near collision leads to one robot stopping, and hence an increase in delay time.) This suggests that the mechanism spreads the robots more evenly throughout the space. The total distance traveled (Figure 5(f)) is greater for greedy-taxi than simple auction, as above, because the bidding strategy employed for the auction mechanism attempts to minimize distance. The results for the simulation also reflect these findings.

Figure 6 provides greater detail. It lists results by scenario—so it is possible to compare specific sets of interest points. It also shows the results of each trial for each scenario, to provide a sense of the distribution of the results. Rectangular boxes enclose the range of values for each measurement; enclosed markers indicate individual values. Figure 6(a) shows that the superior performance of auction over greedy for the physical robots broadly holds for run time across all the scenarios. In other words, the results hold for pairwise comparisons between the dark boxes for Aa versus Ga, Ab versus Gb, and so on. Results on the deliberation time also clearly hold across all scenarios (Figure 6(b)), with the greedy mechanism allocating points more quickly than the auction, as would be expected. The situation with the rest of the results is less clear. For idle time (Figure 6(c)), it seems the main difference between the two allocation mechanisms may be the effect of the large amount of idle time in one trial of the greedy mechanism in scenario c, while scenarios c and d have the biggest effect on the number of near collisions (Figure 6(c)), and hence on delay time (Figure 6(d)).

Figure 6 also provides more detail on the differences between results produced with the physical robots and results produced with the simulation, with an eye to identifying which metrics transfer between the two and might used to predict performance [27]. While the simulation doesn't provide accurate information about the absolute values of runtime or deliberation time (due to the different types of processors and network configurations used in each experimental environment), it does predict the relative relationship between the measured values for different scenarios. It tracks delay time and speed fairly closely—the aim is for the range of values measured in simulation to fall within the range of values measured in the physical environment. The best results are for speed (Figure 6(f)), where all but one of the simulated results falls within the range of those for the physical robots. Given the propensity of the physical robots to generate outliers (through some disastrous temporary mis-localization, for example), this alignment is probably as good as one might reasonably expect.

Figure 7 breaks down some of the metrics by robot. Robots were used consistently across the experiments, that is, what we call $robot_i$ started in the same place for every experiment and interacted with the task allocation mechanism in the same order for every experiment, for all i.[4] Thus these results provide another view of how the allocation mechanisms compare. For distance travelled (Figure 7(a)), travel time (Figure 7(b)) and delay time (Figure 7(d)), in the physical environment, the auction mechanism not only improves the performance of the team, but also improves the performance of each robot. Idle time (Figure 7(c)), on the other hand, is more nuanced. While idle time decreased for three of the robots with the auction rather than greedy allocation in the physical robots, it actually increased for the last robot. Presumably this is because that robot consistently gets easier allocations and so spends more time idle. When the tasks cannot be evenly shared—as in the extreme case where all tasks take equally long and there is one more task than robots—an allocation that is efficient for the team on some metrics may end up leaving some team members sitting idle. Again, the simulated results mirror the results on the physical robots.

[4] The batteries on the Blackfin are not easily removed, so not all the experiments used the same set of robots. Some robots sat out some experiments while their batteries were being recharged. Thus different physical robots played the role of different $robot_i$ in different experiments.

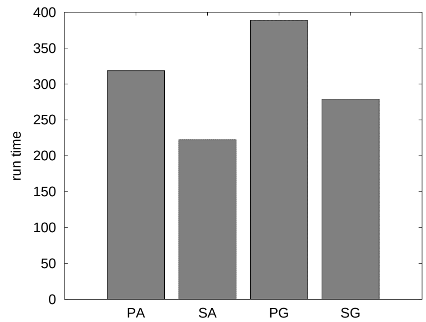

(a) Overall run time

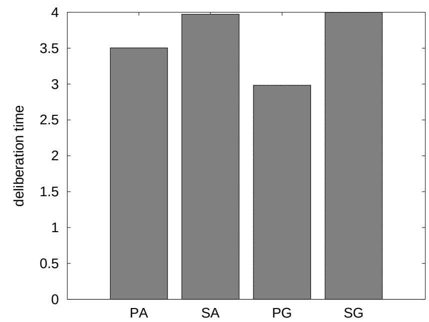

(b) Deliberation time

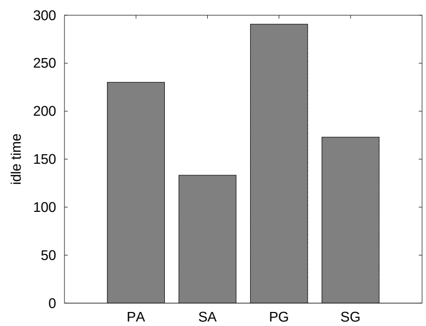

(c) Idle time

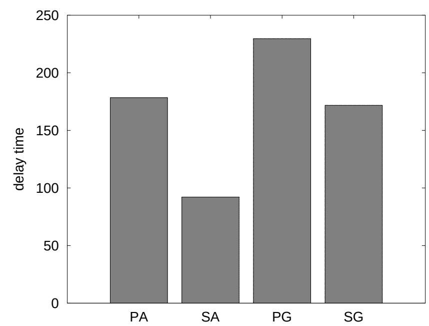

(d) Delay time

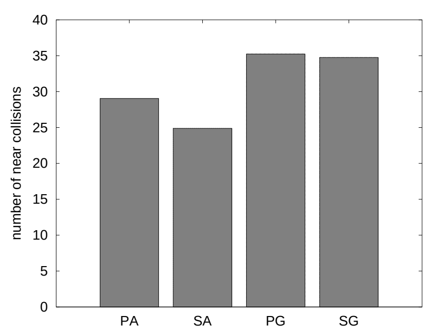

(e) Near collisions

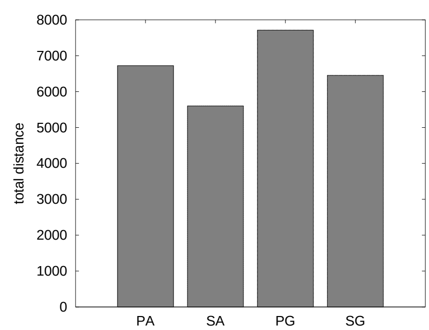

(f) Distance traveled

Fig. 5. Combined results for all scenarios. Values for each metric distinguish between the physical (P) and simulated (S) experiments, and between the simple auction (A) and greedy-taxi (G) task allocation.

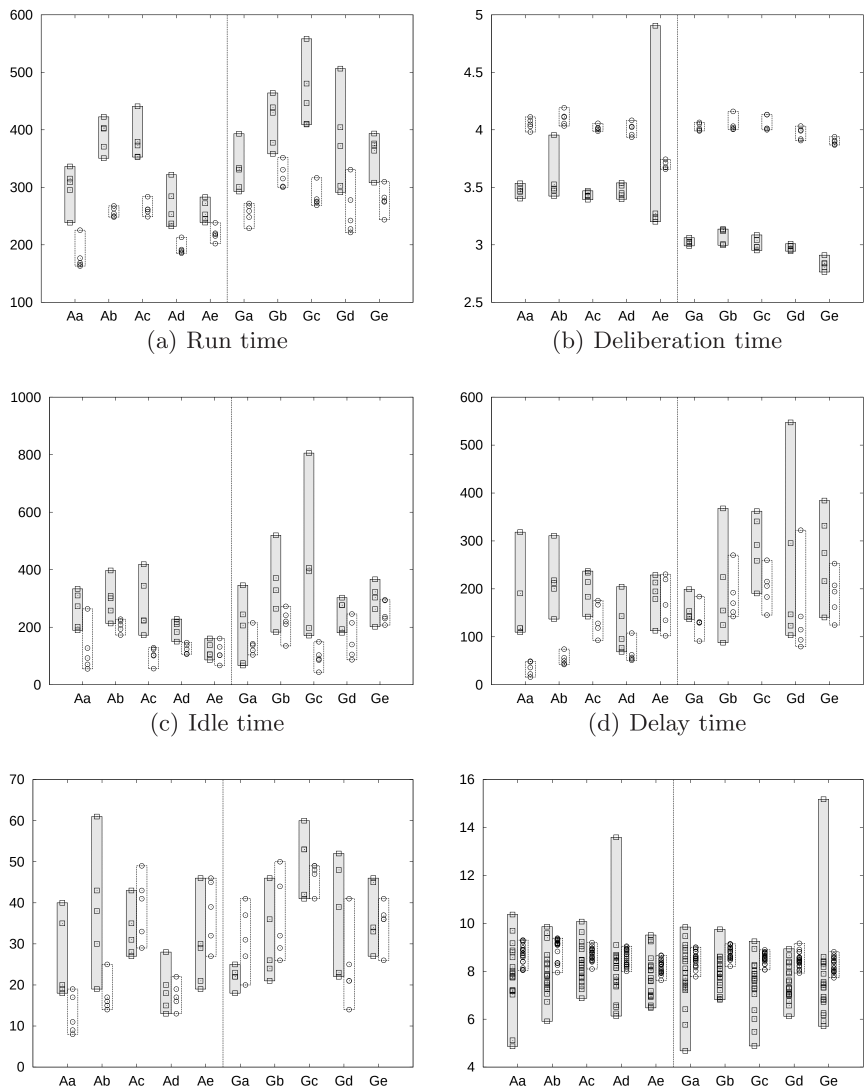

(a) Run time (b) Deliberation time

(c) Idle time (d) Delay time

(e) Near collisions (f) Speed

Fig. 6. Results from individual scenarios. We compare the range of values measured in the physical and simulated environments. Results for the physical robots are enclosed in grey boxes; results for the simulated robots are enclosed in white boxes. We distinguish between simple auction (A) and greedy-taxi (G) task allocation, for each of the five different scenarios, labelled a through e.

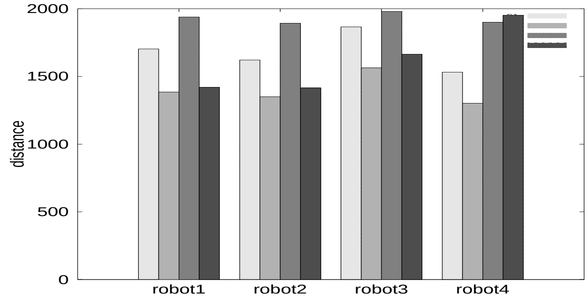

(a) Distance travelled

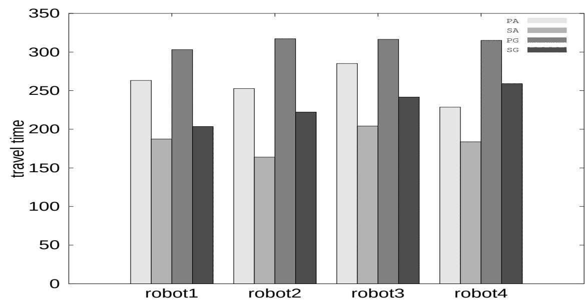

(b) Travel time

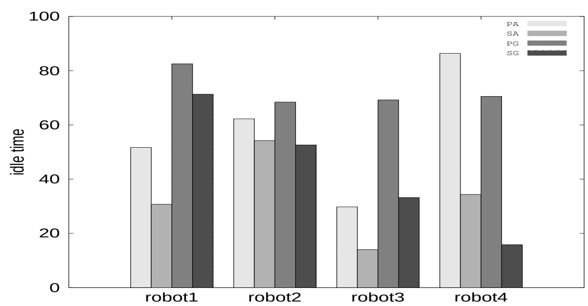

(c) Idle time

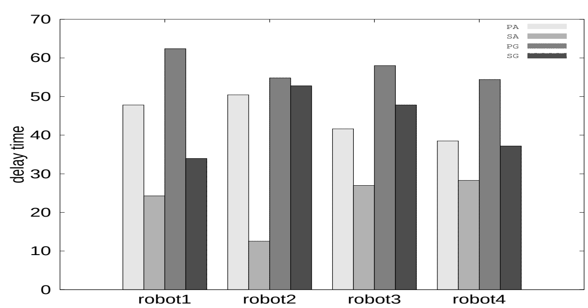

(d) Delay time

Fig. 7. The results for each of the four robots, distinguished by their identifier within the code. The starting position of a given robot determines its identifier. See Figure 3 (top left) for a diagram of the robots' starting positions.

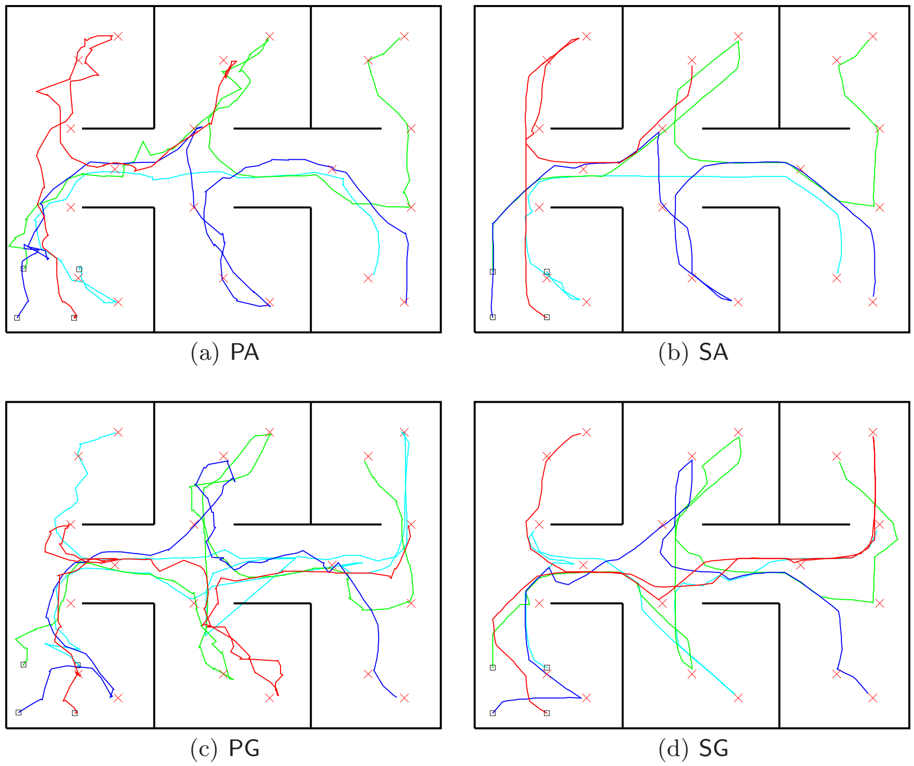

Fig. 8. Robot trajectories in the physical (left) and simulated (right) environments, for scenario a, using task allocation by simple auction (top) and greedy-taxi (bottom) mechanism.

Figure 8 illustrates paths traveled by the robots in the physical and simulated environments for a single scenario (a) using both allocation mechanisms. The top two graphs show the simple auction allocation mechanism and the bottom two graphs show the greedy-taxi allocation mechanism. Comparing the two graphs in the top row, where physical and simulated robots received the same distribution of interest points, it becomes obvious why the physical robots travelled further: the paths are clearly more irregular, while the simulated robots' paths are smoother. This is due to errors in the camera-based localization and noise from the robot's motors that cause the robots to move in jerky motions, especially when turning. Comparing the pairs of graphs in each column, especially for the simulated runs (rightmost column), it is obvious that the robot team using the simple auction allocation mechanism traced shorter and more efficient paths than the team using the greedy-taxi allocation mechanism.

4.4 Discussion

The headline result is that the team can carry out its work more efficiently with the auction mechanism for task allocation than with greedy allocation. This confirms the hypothesis tested here, not a surprising result but gratifying in light of our larger aim. Since, in the long term, we are testing the hypothesis that multiagent coordination mechanisms can be useful in multi-robot teams, it would be distressing were a change from greedy task allocation to an allocation mechanism that considered some information about the task to have no effect on performance. It would also suggest that our longer term hypothesis was false. With this result, we are ready for the next step, to test more complex task allocation mechanisms.

As indicated above, speed is the only metric for which there is consistently good agreement between the absolute values of the simulated and physical results. Nonetheless, there is good qualitative agreement for runtime, deliberation time and delay time; the order across the scenarios is the same for both simulated and physical tests. This suggests that the simulation can predict results such as which scenarios will take longer to allocate and execute. The observant reader may notice that run time is faster in simulation, while deliberation time is faster with the physical robots; though we also point out that this difference is negligible since the differences in deliberation time (in Figure 5(b)) are two orders of magnitude smaller than run time (in Figure 5(a)). Referring to the footnotes in Section 4.1 that describe the details of the processors used in the physical and simulated environments, the (albeit negligible) difference is most likely due to the additional processors employed for the physical experiments.

A final point for discussion is why run time (in Figure 6(a)) is consistently shorter for the simulator than for the physical robots. At first glance, this suggests that the simulator is not well-calibrated to the physical environment, but the reason is more subtle. The distance travelled by the robots (in Figure 7(b)) shows that the simulated robots travel shorter distances. Figure 8 illustrates this point more clearly. The discrepancy arises because in the simulator the robots always know where they are, and they suffer none of the problems of poor localization to which real robots are subject. Even with the overhead camera grid to identify their position, the framework has problems with misidentification, blind spots, and overlapping pictures from two cameras. In addition, the Blackfins are particularly prone to errors in orientation because of their skid steering, which is implicit in their tracked design. The high power required for skid steering [25] makes the robots turn quickly and not very precisely. Since the physical robots travel further, they will naturally take longer to complete their tasks, and the results for speed (Figure 6(f)), which match up closely between physical and simulated robots, strongly suggest that the calibration is correct.

5 Conclusions and Future Work

The intent of this paper was to test the hypothesis that a simple auction mechanism to allocate interest points to robots would provide a better allocation than a first-come first-served greedy-taxi mechanism. The results show this hypothesis to be true—on almost all metrics, the simple auction mechanism is superior, both in the physical and the simulated environments.

This is not a surprising result, but it is a gratifying one from the perspective of our long-term research focus. The improvement over a greedy baseline with a simple coordination mechanism suggests that our rough-and-ready framework is not too rough for this kind of research. It also suggests, because the environment is so rough, that what appeared here will be replicated in different environments and on different tasks.

Our results also suggest a line of future research: to look for task allocation mechanisms that outperform the simple auction. We have suggested that a combinatorial auction would likely do better, and we will test that thesis next. It also seems likely that a series of sequential auctions in which the bid for each subsequent point reflected the total path cost rather than just the cost of the path that must be added would show an improvement, and we will test that as well.

Acknowledgments

This work was supported by the National Science Foundation under #CNS-0851901, #CNS-0520989, #IIS-1116843 and #IIS-1117000. The authors are grateful to Ofear Balas for his improvements to our path planner code. We would also like to thank anonymous reviewers of earlier versions of this paper for helpful comments.

References

1. M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. M. Griffin, and A. Kleywegt. Robot Exploration with Combinatorial Auctions. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Las Vegas*, pages 1957–1962, 2003.
2. M B. Dias and A. Stentz. A Free Market Architecture for Distributed Control of a Multirobot System. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6)*, pages 115–122, July 2000.
3. M. B. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: A survey and analysis. Technical Report CMU-RI-TR-05-13, Carnegie Mellon University, 2005.
4. S. L. Epstein. FOr the Right Reasons: The FORR architecture for learning in a skill domain. *Cognitive Science*, 18(3), 1994.
5. S. L. Epstein, E. Schneider, A. T. Özgelen, J. P. Mũnoz, M. Costantino, S. Parsons, and E. I. Sklar. Applying FORR to human/multi-robot teams. In *Workshop on Human-Agent-Robot Teams (HART) at the Human-Robot Interaction (HRI) conference*, Boston, MA, USA, March 2012.
6. V. Frias-Martinez, E. I. Sklar, and S. Parsons. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proceedings of the 8th RoboCup International Symposium*, 2004.

7. Y. Fujishim, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, Stockholm, 1999.
8. B. Gerkey, R. T. Vaughan, and A. Howard. The Player/Stage Project: Tools for Multi-Robot and Distributed Sensor Systems. In *Proceedings of the 11th International Conference on Advanced Robotics (ICAR)*, 2003.
9. B. P. Gerkey and M. J Mataríc. Sold!: Auction methods for multi-robot control. *IEEE Transactions on Robotics and Automation Special Issue on Multi-Robot Systems*, 18(5), 2002.
10. M. Golfarelli, D. Maio, and S. Rizzi. A Task-Swap Negotiation Protocol Based on the Contract Net Paradigm. Technical Report 005-97, DEIS, CSITE - Università di Bologna, 1997.
11. M. K. Habib. Humanitarian Demining: Reality and the Challenge of Technology. *International Journal of Advanced Robotic Systems*, 4(2), 2007.
12. A. Jacoff, E. Messina, and J. Evans. A standard test course for urban search and rescue robots. In *Proceedings of the Workshop on Performance Metrics for Intelligent Systems (PerMIS)*, 2000.
13. N. Kalra and A. Stenz. A Market Approach to Tightly-Coupled Multi-Robot Coordination: First Results. In *Proceedings of the ARL Collaborative Technologies Alliance Symposium*, 2003.
14. S. Koenig, P. Keskinocak, and C. Tovey. Progress on agent coordination with cooperative auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2010.
15. 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 *Proceedings of National Conference on Artificial Intelligence (AAAI)*, pages 1625–1629, 2006.
16. S. Kraus. Automated negotiation and decision making in multiagent environments. In *In: MultiAgent Systems and Applications. ACAI-EASSS 2001 Proceedings, Luck M., Marik V., Stepankova O., Trappl R. (eds). Springer-Verlag*, pages 150–172. Springer-Verlag, 2001.
17. M. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. Kelywegt. Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the International Conference on Intelligent Robotics and Systems (IROS)*, 2004.
18. M. Lagoudakis, V. Markakis, D. Kempe, P. Keskinocak, S. Koenig, A. Kleywegt, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Proceedings of Robotics: Science and Systems Conference*, 2005.
19. M. Mataric, G. Sukhatme, and E. Ostergaard. Multi-robot task allocation in uncertain environments. *Autonomous Robots*, 2003.
20. N. Michael, M. M. Zavlanos, V. Kumar, and G. J. Pappas. Distributed multi-robot task assignment and formation control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2008.
21. R. R. Murphy, J. Casper, and M. Micire. Potential tasks and research issues for mobile robots in robocup rescue. In *Robot Soccer World Cup IV, LNAI 2019.* Springer Verlag, 2001.
22. T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, 2002.
23. P. F. Santana, J. Barata, and L. Correia. Sustainable Robots for Humanitarian Demining. *International Journal of Advanced Robotic Systems*, 4(2), 2007.

24. S. Sariel and T. Balch. Efficient bids on task allocation for multi-robot exploration. In *Proceedings of the Nineteenth Florida Artificial Intelligence Research Society Conference*, 2006.
25. B. Shamah. Experimental comparison of skid steering vs. explicit steering for a wheeled mobile robot. Master's thesis, The Robotics Institute, Carnegie Mellon University, Pittsburgh Pennsylvania 15213, March 1999. CMU-RI-TR-99-06.
26. E. I. Sklar, A. T. Özgelen, J. P. Muñoz, J. Gonzalez, M. Manashirov, S. L. Epstein, and S. Parsons. Designing the HRTeam framework: Lessons learned from a rough-and-ready human/multi-robot team. In *Proceedings of the Workshop on Autonomous Robots and Multirobot Systems*, Taipei, Taiwan, May 2011.
27. E. I. Sklar, A. T. Özgelen, E. Schneider, M. Costantino, J. P. Muñoz, S. L. Epstein, and S. Parsons. On transfer from multiagent to multi-robot systems. In *Proceedings of the Workshop on Autonomous Robots and Multirobot Systems*, Valencia, Spain, June 2012.
28. R. G. Smith. The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. In Alan H. Bond and Les Gasser, editors, *Distributed Artificial Intelligence*, pages 357–366. Morgan Kaufmann Publishers Inc., 1988.
29. C. Tovey, M. G. Lagoudakis, S. Jain, and S. Koenig. Generation of Bidding Rules for Auction-Based Robot Coordination. In *Proceedings of the 3rd International Multi-Robot Systems Workshop, Naval Research Laboratory, Washington DC*, March 2005.
30. M. P. Wellman and P. R. Wurman. Market-Aware Agents for a Multiagent World. *Robotics and Autonomous Systems*, 24:115–125, 1998.
31. R. Zlot, A. Stentz, M. B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proceedings of the IEEE Conference on Robotics and Automation*, 2002.