THE GENERATION OF BIDDING RULES FOR AUCTION-BASED ROBOT COORDINATION*

Craig Tovey and Michail G. Lagoudakis
School of Industrial and Systems Engineering, Georgia Institute of Technology
{ctovey, michail.lagoudakis}@isye.gatech.edu

Sonal Jain and Sven Koenig
Computer Science Department, University of Southern California
{sonaljai, skoenig}@usc.edu

Abstract Robotics researchers have used auction-based coordination systems for robot teams because of their robustness and efficiency. However, there is no research into systematic methods for deriving appropriate bidding rules for given team objectives. In this paper, we propose the first such method and demonstrate it by deriving bidding rules for three possible team objectives of a multi-robot exploration task. We demonstrate experimentally that the resulting bidding rules indeed exhibit good performance for their respective team objectives and compare favorably to the optimal performance. Our research thus allows the designers of auction-based coordination systems to focus on developing appropriate team objectives, for which good bidding rules can then be derived automatically.

Keywords: Auctions, Bidding Rules, Multi-Robot Coordination, Exploration.

1. Introduction

The time required to reach other planets makes planetary surface exploration missions prime targets for automation. Sending rovers to other planets either instead of or together with people can also significantly reduce the danger and cost involved. Teams of rovers are both more fault tolerant (through redundancy) and more efficient (through parallelism)

*We thank Apurva Mudgal for his help. This research was partly supported by NSF awards under contracts ITR/AP0113881, IIS-0098807, and IIS-0350584. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, companies or the U.S. government.

than single rovers if the rovers are coordinated well. However, rovers cannot be easily tele-operated since this requires a large number of human operators and is communication intensive, error prone, and slow. Neither can they be fully preprogrammed since their activities depend on their discoveries. Thus, one needs to endow them with the capability to coordinate autonomously with each other. Consider, for example, a multi-robot exploration task where a team of lunar rovers has to visit a number of given target locations to collect rock samples. Each target must be visited by at least one rover. The rovers first allocate the targets to themselves, and each rover then visits the targets that are allocated to it. The rovers know their current location at all times but might initially not know where obstacles are in the terrain. It can therefore be beneficial for the rovers to re-allocate the targets to themselves as they discover more about the terrain during execution, for example, when a rover discovers that it is separated by a big crater from its next target. Similar multi-robot exploration tasks arise for mine sweeping, search and rescue operations, police operations, and hazardous material cleaning, among others.

Multi-robot coordination tasks are typically solved with heuristic methods since optimizing the performance is often computationally intractable. They are often solved with decentralized methods since centralized methods lack robustness: if the central controller fails, so does the entire robot team. Market mechanisms, such as auctions, are popular decentralized and heuristic multi-robot coordination methods (Rabideau et al., 2000). In this case, the robots are the bidders and the targets are the goods up for auction. Every robot bids on targets and then visits all targets that it wins. As the robots discover more about the terrain during execution, they run additional auctions to change the allocation of targets to themselves. The resulting auction-based coordination system is efficient in terms of communication (robots communicate only numeric bids) and computation (robots compute their bids in parallel). It is therefore not surprising that auctions have been shown to be effective multi-robot coordination methods (Gerkey and Matarić, 2002; Zlot et al., 2002; Thayer et al., 2000; Goldberg et al., 2003). However, there are currently no systematic methods for deriving appropriate bidding rules for given team objectives. In this paper, we propose the first such method and demonstrate it by deriving bidding rules for three possible team objectives of the multi-robot exploration task. We demonstrate experimentally that the resulting bidding rules indeed exhibit good performance for their respective team objectives and compare favorably to the optimal performance. Our research thus allows the designers of auction-based coordination systems to focus on developing appropriate team objectives, for which good bidding rules can then be derived automatically.

2. The Auction-Based Coordination System

In known environments, all targets are initially unallocated. During each round of bidding, all robots bids on all unallocated targets. The robot that places the overall lowest bid on any target is allocated that particular target. A new round of bidding starts, and all robots bid again on all unallocated targets, and so on until all targets have been allocated to robots. (Note that each robot needs to bid only on a single target during each round, namely on one of the targets for which its bid is the lowest, since all other bids from the same robot have no chance of winning.) Each robot then calculates the optimal path for the given team objective for visiting the targets allocated to it and then moves along that path. A robot does not move if no targets are allocated to it.

In unknown environments, the robots proceed in the same way but under the optimistic initial assumption that there are no obstacles. As the robots move along their paths and a robot discovers a new obstacle, it informs the other robots about it. Each robot then re-calculates the optimal path for the given team objective for visiting the unvisited targets allocated to it, taking into account all obstacles that it knows about. If the performance significantly degrades for at least one robot (in our experiments, we use a threshold of 10 percent difference), then the robots use auctions to re-allocate all unvisited targets among themselves. Each robot then calculates the optimal path for the given team objective for visiting the targets allocated to it and then moves along that path, and so on until all targets have been visited.

This auction-based coordination system is similar to multi-round auctions and sequential single-item auctions. Its main advantage is its simplicity and the fact that it allows for a decentralized implementation on real robots. Each robot computes its one bid locally and in parallel with the other robots, broadcasts the bid to the other robots, listens to the broadcasts of the other robots, and then locally determines the winning bid. Thus, there is no need for a central auctioneer and therefore no single point of failure. A similar but more restricted auction scheme has been used in the past for robot coordination (Dias and Stentz, 2000).

3. Team Objectives for Multi-Robot Exploration

A multi-robot exploration task consists of the locations of n robots and m targets as well as a cost function that specifies the cost of moving between locations. The objective of the multi-robot exploration task is to find an allocation of targets to robots and a path for each robot that visits all targets allocated to it so that the team objective is achieved. Note that the robots are not required to return to their initial locations. In this paper, we study three team objectives:

- MiniSum: Minimize the sum of the path costs over all robots.
- MiniMax: Minimize the maximum path cost over all robots.
- MiniAve: Minimize the average per target cost over all targets.

The path cost of a robot is the sum of the costs along its path, from its initial location to the first target on the path, and so on, stopping at the last target on the path. The per target cost of a target is the sum of the costs along the path of the robot that visits the target in question, from its initial location to the first target on the path, and so on, stopping at the target in question.

Optimizing the performance for the three team objectives is NP-hard and thus likely computationally intractable, as they resemble the Traveling Salesperson Problem, the Min-Max Vehicle Routing Problem, and the Traveling Repairperson Problem (or Minimum Latency Problem), respectively, which are intractable even on the Euclidean plane. However, these team objectives cover a wide range of applications. For example, if the cost is energy consumption, then the MiniSum team objective minimizes the total energy consumed by all robots until all targets have been visited. If the cost is travel time, then the Minimax team objective minimizes the time until all targets have been visited (task-completion time) and the MiniAve team objective minimizes how long it takes on average until a target is visited (target-visit time). The MiniSum and MiniMax team objectives have been used in the context of multi-robot exploration (Dias and Stentz, 2000; Dias and Stentz, 2002; Berhault et al., 2003; Lagoudakis et al., 2004). The MiniAve team objective, on the other hand, has not been used before in this context although it is very appropriate for search-and-rescue tasks, where the health condition of several victims deteriorates until a robot visits them. Consider, for example, an earthquake scenario where an accident site with one victim is located at a travel time of 20 units to the west of a robot and another accident site with twenty victims is located at a travel time of 25 units to its east. In this case, visiting the site to the west first and then the site to the east achieves both the MiniSum and the MiniMax team objectives. However, the twenty victims to the east are visited very late and their health condition thus is very bad. On the other hand, visiting the site to the east first and then the site to the west achieves the MiniAve team objective and results in an overall better average health condition of the victims. This example illustrates the importance of the MiniAve team objective in cases where the targets occur in clusters of different sizes.

4. Systematic Generation of Bidding Rules

We seek to derive an appropriate bidding rule for a given team objective. This problem has not been studied before in the robotics literature. Assume that there are n robots $r_1, \ldots, r_n$ and m currently unallocated targets $t_1, \ldots, t_m$. Assume further that the team objective has the structure to assign a set of targets T_i to robot r_i for all i, where the sets $T = \{T_1, \ldots, T_n\}$ form a partition of all targets that optimizes the performance $f\big(g(r_1, T_1), \ldots, g(r_n, T_n)\big)$ for given functions f and g. Function g determines the performance of each robot, and function f determines the performance of the team as a function of the performance of the robots. The three team objectives fit this structure. For any robot r_i and any set of targets T_i, let $PC(r_i, T_i)$ denote the minimum path cost of robot r_i and $STC(r_i, T_i)$ denote the minimum sum of per target costs over all targets in T_i if robot r_i visits all targets in T_i from its current location. Then, it holds that

- MiniSum: $\min_T \sum_j PC(r_j, T_j)$,
- MiniMax: $\min_T \max_j PC(r_j, T_j)$, and
- MiniAve: $\min_T \frac{1}{m} \sum_j STC(r_j, T_j)$.

A bidding rule determines how much a robot bids on a target. We propose the following bidding rule for a given team objective, which is directly derived from the team objective itself.

Bidding Rule Robot r bids on target t the difference in performance for the given team objective between the current allocation of targets to robots and the allocation that results from the current one if robot r is allocated target t. (Unallocated targets are ignored.)

Consequently, robot r_i should bid on target t

$$f\big(g(r_1, T_1'), \ldots, g(r_n, T_n')\big) - f\big(g(r_1, T_1), \ldots, g(r_n, T_n)\big),$$

where $T_i' = T_i \cup \{t\}$ and $T_j' = T_j$ for $i \neq j$. The bidding rule thus performs hill climbing to maximize the performance and can thus suffer from local optima. However, optimizing the performance is NP-hard for the three team objectives. Our auction-based coordination system is therefore not designed to optimize the performance but to be efficient and result in a good performance, and hill climbing has these properties. One potential problem with the bidding rule is that the robots might not have all the information needed to compute the bids. For example, a robot may not know the locations of the other robots. However, we will now show that a robot can calculate its bids for the three team objectives knowing only its current location, the set of targets allocated to it, and the cost function:

- For the MiniSum team objective, robot r_i should bid on target t

$$\sum_j PC(r_j, T_j') - \sum_j PC(r_j, T_j) = PC(r_i, T_i \cup \{t\}) - PC(r_i, T_i).$$

- For the MiniMax team objective, robot r_i should bid on target t

$$\max_j PC(r_j, T_j') - \max_j PC(r_j, T_j) = PC(r_i, T_i \cup \{t\}) - \max_j PC(r_j, T_j).$$

  This derivation uses the fact that $\max_j PC(r_j, T_j') = PC(r_i, T_i')$, otherwise target t would have already been allocated in a previous round of bidding. The term $\max_j PC(r_j, T_j)$ can be dropped since the outcomes of the auctions remain unchanged if all bids change by a constant. Thus, robot r_i can bid just $PC(r_i, T_i \cup \{t\})$ on target t.

- For the MiniAve team objective, robot r_i should bid on target t

$$\frac{1}{m}\sum_j STC(r_j, T_j') - \frac{1}{m}\sum_j STC(r_j, T_j) = \frac{1}{m}\Big(STC(r_i, T_i \cup \{t\}) - STC(r_i, T_i)\Big).$$

  The factor $1/m$ can be dropped since the outcomes of the auctions remain unchanged if all bids are multiplied by a constant factor. Thus, robot r_i can bid just $STC(r_i, T_i \cup \{t\}) - STC(r_i, T_i)$ on target t.

Thus, the bidding rules for the three team objectives are

- BidSum: $PC(r_i, T_i \cup \{t\}) - PC(r_i, T_i)$,
- BidMax: $PC(r_i, T_i \cup \{t\})$, and
- BidAve: $STC(r_i, T_i \cup \{t\}) - STC(r_i, T_i)$.

The robots need to be able to calculate their bids efficiently but computing $PC(r_i, T_i \cup \{t\})$ or $STC(r_i, T_i \cup \{t\})$ is NP-hard. Robot r_i thus uses a greedy method to approximate these values. In particular, it finds a good path that visits the targets in $T_i \cup \{t\}$ for a given team objective as follows. It already has a good path that visits the targets in T_i.

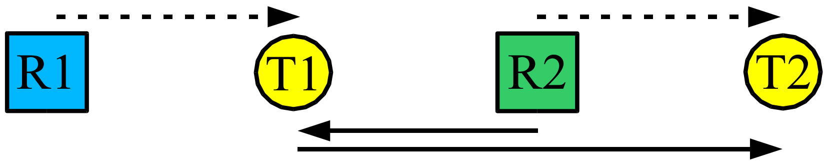

Figure 1. A simple multi-robot exploration task.

First, it inserts target t into all positions on the existing path, one after the other. Then, it tries to improve each new path by first using the 2-opt improvement rule and then the 1-target 3-opt improvement rule. Finally, it picks the best one of the resulting paths for the given team objective. The 2-opt improvement rule takes a path and inverts the order of targets in each one of its continuous subpaths in turn, picks the best one of the resulting paths for the given team objective, and repeats the procedure until the path can no longer be improved. The 1-target 3-opt improvement rule removes a target from the path and inserts it into all other possible positions on the path, picks the best one of the resulting paths for the given team objective, and repeats the procedure until the path can no longer be improved.

The three bidding rules are not guaranteed to achieve their respective team objectives even if the values $PC(r_i, T_i \cup \{t\})$ and $STC(r_i, T_i \cup \{t\})$ are computed exactly. Consider the simple multi-robot exploration task in Figure 1 with 2 robots and 2 targets and unit costs between adjacent locations. All bidding rules can result in the robots following the solid lines, resulting in a performance of 3 for the MiniSum team objective, a performance of 3 for the MiniMax team objective, and a performance of 2 for the MiniAve team objective. However, the robots should follow the dashed lines to maximize the performance for all three team objectives, resulting in a performance of 2 for the MiniSum team objective, a performance of 1 for the MiniMax team objective, and a performance of 1 for the MiniAve team objective. (We rely on a particular way of breaking ties in this multi-robot exploration example but can easily change the edge costs by small amounts to guarantee that the bidding rules result in the robots following the solid lines independently of how ties are broken.) In a forthcoming paper, we analyze the performance of the three bidding rules theoretically and show that the performance of the BidSum bidding rule in the Euclidean case is at most a factor of two away from optimum, whereas no constant-factor bound exists for the performance of the BidMax and BidAve bidding rules even in the Euclidean case.

5. Experimental Evaluation

To demonstrate that the performance of the three bidding rules is indeed good for their respective team objectives, we implemented them and then tested them in office-like environments with rooms, doors, and corridors, as shown in Figure 2. We performed experiments with both unclustered and clustered targets. The locations of the robots and targets for each multi-robot exploration task were chosen randomly in the unclustered target case. The locations of the robots and targets were also chosen randomly in the clustered target case, but with the restriction that 50 percent of the targets were placed in clusters of 5 targets each. The numbers in the tables below are averages over 10 different multi-robot exploration tasks with the same settings. The performance of the best bidding rule for a given team objective is shown in bold.

5.1 Known Environments

We mapped our environments onto eight-connected uniform grids of size 51×51 and computed all costs between locations as the shortest distances on the grid. Our auction-based coordination system used these costs to find an allocation of targets to robots and a path for each robot that visits all targets allocated to it. We interfaced it to the popular Player/Stage robot simulator (Gerkey et al., 2003) to execute the paths and visualize the resulting robot trails. Figure 2 shows the initial locations of the robots (squares) and targets (circles) as well as the resulting robot trails (dots) for each one of the three bidding rules for a sample multi-robot exploration task with 3 robots and 20 unclustered targets in a completely known environment. SUM, MAX and AVE in the caption of the figure denote the performance for the MINISUM, MINIMAX and MINIAVE team objectives, respectively. Each bidding rule results in a better performance for its team objective than the other two bidding rules. For example, the BIDSUM bidding rule results in paths of very different lengths, whereas the BIDMAX bidding rule results in paths of similar lengths. Therefore, the performance of the BIDMAX bidding rule is better for the MINIMAX team objective than the one of the BIDSUM bidding rule.

We compared the performance of the three bidding rules against the optimal performance for multi-robot exploration tasks with one or two robots and ten targets. The optimal performance was calculated by formulating the multi-robot exploration tasks as integer programs and solving them with the commercial mixed integer program solver CPLEX. The NP-hardness of optimizing the performance did not allow us to solve larger multi-robot exploration tasks. Table 1 shows the performance of

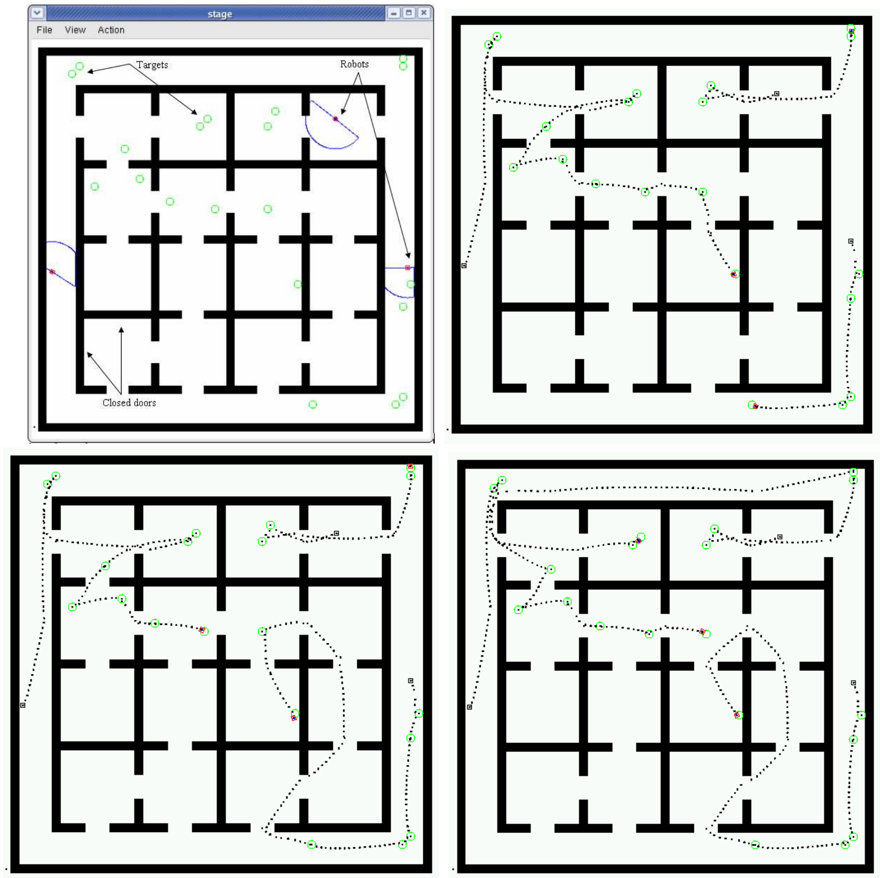

Figure 2. Player/Stage screenshots: initial locations (top left) and robot trails with the BidSum (top right) [Sum=182.50, Max=113.36 , Ave=48.61], BidMax (bottom left) [Sum=218.12 , Max=93.87 , Ave=46.01], and BidAve (bottom right) [Sum=269.27 , Max=109.39 , Ave=45.15] bidding rules.

each bidding rule and the optimal performance for each team objective. Again, each bidding rule results in a better performance for its team objective than the other two bidding rules, with the exception of ties between the BidSum and BidMax bidding rules for multi-robot exploration tasks with one robot. These ties are unavoidable because the MiniSum and MiniMax team objectives are identical for one-robot exploration tasks. The performance of the best bidding rule for each team objective is always close to the optimal performance. In particular, the performance of the BidSum bidding rule for the MiniSum team objective is within a factor of 1.10 of optimal, the performance of the BidMax bidding rule for the MiniMax team objective is within a factor

Table 1. Performance of bidding rules against optimal in known environments.

Robots	Bidding Rule	Unclustered			Clustered		
		SUM	MAX	AVE	SUM	MAX	AVE
1	BIDSUM	**199.95**	**199.95**	103.08	**143.69**	**143.69**	78.65
1	BIDMAX	**199.95**	**199.95**	103.08	**143.69**	**143.69**	78.65
1	BIDAVE	214.93	214.93	**98.66**	155.50	155.50	**63.12**
1	OPTIMAL	199.95	199.95	98.37	143.69	143.69	63.12
2	BIDSUM	**193.50**	168.50	79.21	**134.18**	97.17	62.47
2	BIDMAX	219.15	**125.84**	61.39	144.84	**90.10**	57.38
2	BIDAVE	219.16	128.45	**59.12**	157.29	100.56	**49.15**
2	OPTIMAL	189.15	109.34	55.45	132.06	85.86	47.63

of 1.44 of optimal, and the performance of the BIDAVE bidding rule for the MINIAVE team objective is within a factor of 1.28 of optimal.

We also compared the performance of the three bidding rules against each other for large multi-robot exploration tasks with one, five or ten robots and 100 targets. Table 2 shows the performance of each bidding rule. Again, each bidding rule results in a better performance for its team objective than the other two bidding rules, with the exception of the unavoidable ties.

5.2 Unknown Environments

We compared the performance of the three bidding rules against each other for the same large multi-robot exploration tasks as in the previous section but in initially completely unknown environments. In this case, we mapped our environments onto four-connected uniform grids of size 51×51 and computed all costs between locations as the shortest distances on the grid. These grids were also used to simulate the movement of the robots in a coarse and noise-free simulation. (We could not use eight-connected grids because diagonal movements are longer than horizontal and vertical ones, and the simulation steps thus would need to be much smaller than moving from cell to cell.) The robots sense all blockages in their immediate four-cell neighborhood. Table 3 shows the performance of each bidding rule. Again, each bidding rule results in a better performance for its team objective than the other two bidding rules, with the exception of the unavoidable ties and two other exceptions. The average number of auctions is 28.37 with a maximum of 82 auctions in one case. In general, the number of auctions increases with the number of robots. Note that the difference in performance between known and unknown environments is at most a factor of three. It is remarkable that our auction-based coordination system manages to achieve such a good

Table 2. Performance of bidding rules against each other in known environments.

Robots	Bidding Rule	Unclustered			Clustered		
		SUM	MAX	AVE	SUM	MAX	AVE
1	BIDSUM	**554.40**	**554.40**	281.11	**437.25**	**437.25**	212.81
1	BIDMAX	**554.40**	**554.40**	281.11	**437.25**	**437.25**	212.81
1	BIDAVE	611.50	611.50	**243.30**	532.46	532.46	**169.20**
5	BIDSUM	**483.89**	210.30	80.74	**374.33**	186.50	66.94
5	BIDMAX	548.40	**130.41**	58.70	450.72	**112.18**	50.50
5	BIDAVE	601.28	146.18	**55.19**	500.05	132.98	**42.41**
10	BIDSUM	**435.30**	136.70	45.89	**318.52**	102.15	35.14
10	BIDMAX	536.90	**77.95**	31.39	402.30	**63.89**	25.88
10	BIDAVE	564.73	88.23	**30.04**	437.23	71.52	**22.02**

Table 3. Performance of bidding rules against each other in unknown environments.

Robots	Bidding Rule	Unclustered			Clustered		
		SUM	MAX	AVE	SUM	MAX	AVE
1	BIDSUM	**1459.90**	**1459.90**	**813.40**	**1139.20**	**1139.20**	672.14
1	BIDMAX	**1459.90**	**1459.90**	**813.40**	**1139.20**	**1139.20**	672.14
1	BIDAVE	1588.50	1588.50	826.82	1164.40	1164.40	**463.14**
5	BIDSUM	**943.60**	586.90	223.47	**771.40**	432.90	166.60
5	BIDMAX	979.00	**238.10**	98.48	811.30	216.90	86.58
5	BIDAVE	992.10	240.10	**90.54**	838.30	**214.10**	**79.36**
10	BIDSUM	**799.50**	312.20	93.69	**596.10**	223.20	63.95
10	BIDMAX	885.40	**123.60**	48.43	677.80	**110.60**	37.92
10	BIDAVE	871.80	133.00	**45.19**	697.80	121.50	**35.43**

performance for all team objectives since there has to be some performance degradation given that we switched both from known to unknown environments and from eight-connected to four-connected grids.

6. Conclusions and Future Work

In this paper, we described an auction-based coordination system and then proposed a systematic method for deriving appropriate bidding rules for given team objectives. We then demonstrated it by deriving bidding rules for three possible team objectives of a multi-robot exploration task, that relate to minimizing the total energy consumption, task-completion time, and average target-visit time. (The last team objective had not been used before but we showed it to be appropriate for search-and-rescue tasks.) Finally, we demonstrated experimentally that the derived bidding rules indeed exhibit good performance for their respective team objectives and compare favorably to the optimal performance. In the future, we intend to adapt our methodology to other multi-robot coordination tasks. For example, we intend to study multi-robot coordination with auction-based coordination systems in the presence of additional constraints, such as compatibility constraints which dictate that certain targets can only be visited by certain robots.

References

Berhault, M., Huang, H., Keskinocak, P., Koenig, S., Elmaghraby, W., Griffin, P., and Kleywegt, A. (2003). Robot exploration with combinatorial auctions. In *Proceedings of the International Conference on Intelligent Robots and Systems*, pages 1957–1962.

Dias, M. and Stentz, A. (2000). A free market architecture for distributed control of a multirobot system. In *Proceedings of the International Conference on Intelligent Autonomous Systems*, pages 115–122.

Dias, M. and Stentz, A. (2002). Enhanced negotiation and opportunistic optimization for market-based multirobot coordination. Technical Report CMU-RI-TR-02-18, Robotics Institute, Carnegie Mellon University, Pittsburgh (Pennsylvania).

Gerkey, B. and Matarić, M. (2002). Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768.

Gerkey, B., Vaughan, R., Stoy, K., Howard, A., Sukhatme, G., and Matarić, M. (2003). Most valuable player: A robot device server for distributed control. In *Proceedings of the International Conference on Intelligent Robots and Systems*, pages 1226–1231.

Goldberg, D., Circirello, V., Dias, M., Simmons, R., Smith, S., and Stentz, A. (2003). Market-based multi-robot planning in a distributed layered architecture. In *Proceedings from the International Workshop on Multi-Robot Systems*, pages 27–38.

Lagoudakis, M., Berhault, M., Keskinocak, P., Koenig, S., and Kleywegt, A. (2004). Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the International Conference on Intelligent Robots and Systems*.

Rabideau, G., Estlin, T., Chien, S., and Barrett, A. (2000). A comparison of coordinated planning methods for cooperating rovers. In *Proceedings of the International Conference on Autonomous Agents*, pages 100–101.

Thayer, S., Digney, B., Dias, M., Stentz, A., Nabbe, B., and Hebert, M. (2000). Distributed robotic mapping of extreme environments. In *Proceedings of SPIE: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII*, volume 4195, pages 84–95.

Zlot, R., Stentz, A., Dias, M., and Thayer, S. (2002). Multi-robot exploration controlled by a market economy. In *Proceedings of the International Conference on Robotics and Automation*, pages 3016–3023.