

The Power of Sequential Single-Item Auctions for Agent Coordination

S. Koenig¹ C. Tovey⁴ M. Lagoudakis² V. Markakis³ D. Kempe¹
P. Keskinocak⁴ A. Kleywegt⁴ A. Meyerson⁵ S. Jain⁶

¹ University of Southern California

Computer Science Department
Los Angeles, USA

{skoening, dkempe}@usc.edu

² Technical University of Crete

Department of Electronic and Computer Engineering
Chania, Greece

lagoudakis@intelligence.tuc.gr

³ University of Toronto

Department of Computer Science
Toronto, Canada

vangelis@cs.toronto.edu

⁴ Georgia Institute of Technology

School of Industrial and Systems Engineering
Atlanta, USA

{ctovey, pinar, anton}@isye.gatech.edu

⁵ University of California at Los Angeles

Computer Science Department
Los Angeles, USA

awm@cs.ucla.edu

⁶ Microsoft

US-Windows Client Platform
Redmond, USA

sonalja@microsoft.com

Abstract

Teams of robots are more fault tolerant than single robots, and auctions appear to be promising means for coordinating them. In a recent paper at “Robotics: Science and Systems 2005,” we analyzed a coordination system based on sequential single-item auctions. We showed that the coordination system is simple to implement and computation and communication efficient, and that the resulting sum of all travel distances in known terrain is guaranteed to be only a constant factor away from optimum. In this paper, we put these results in perspective by comparing our coordination system against those based on either parallel single-item auctions or combinatorial auctions, demonstrating that it combines the advantages of both.

Introduction

Consider a team of mobile robots that has to visit a number of given targets (locations) in initially partially unknown terrain. Such exploration tasks are important for environmental clean-up missions, space-exploration missions, and search and rescue missions. It can be necessary or beneficial to re-allocate targets to robots as the robots discover more about the terrain, for example, when a robot discovers that it is separated by a wall from its target. These exploration tasks are special cases of multi-robot task allocation problems (Gerkey & Mataric 2003) and similar to vehicle routing tasks in the presence of changing traffic jams (Fischer *et al.* 1995) but without any hard constraints. To allocate and re-allocate the targets among themselves, the robots can use auctions where they sell and buy targets. The resulting coordination systems are efficient in terms of the required amount of communication since the robots compress their information into numeric bids, and in terms of the required amount of computation since the robots compute their bids in parallel. Auctions have been studied in artificial intelligence starting with the contract net protocol (Smith 1980)

and have subsequently been discussed in the context of robot tasks such as Robosoccer (Nair *et al.* 2002), box pushing (Gerkey & Mataric 2002), security (Kalra, Stentz, & Ferguson 2004) and mapping (Simmons *et al.* 2000). They have also been used on actual robots (Gerkey & Mataric 2002; Zlot *et al.* 2002; Thayer *et al.* 2000). Research in artificial intelligence on coordination systems based on auctions has so far been of an experimental nature. Research in economics has produced analytical results for the allocation of goods with auctions, but has been primarily concerned with issues such as strategic behavior, incentive compatibility, privacy and information. Such issues do not arise in coordination systems based on auctions because one controls both the auction mechanism and the preference structures of the robots. Our research program is intended to close the existing gap by developing a firm theoretical foundation for coordination systems based on auctions. In a recent (highly theoretical) paper at “Robotics: Science and Systems 2005” (Lagoudakis *et al.* 2005), we analyzed a coordination system based on sequential (or multi-round) single-item auctions. We showed that the coordination system is simple to implement and computation and communication efficient, and that the resulting sum of all travel distances in known terrain is guaranteed to be only a constant factor away from optimum. To our knowledge these are the first analytical results about the team performance resulting from using auctions for agent coordination. In this paper, we make these results easily accessible and put them in perspective by comparing against coordination systems that are based on either parallel single-item auctions or combinatorial auctions, demonstrating that their advantages can be combined, which should be of interest to researchers in multi-agent systems, auctions and robotics.

Exploration Task and Approach

We study exploration tasks where a team of robots has to visit a number of targets with known coordinates. Each target has to be visited by at least one robot but the robots do not need to return to their initial locations. The robots

always know their current locations but might initially not know where the obstacles are in the terrain. We study coordination systems for these exploration tasks that fit the following framework: Every robot always follows a minimum-cost path that visits all of the unvisited targets that are allocated to it. Whenever a robot gains more information about the terrain, it shares this information with the other robots. If the remaining path of at least one robot is blocked, then all robots put their unvisited targets up for auction. Each robot then bids in light of the new terrain information. The auction(s) close after a predetermined amount of time and the winning robots get allocated the corresponding targets, which they now own. Then, the cycle repeats. The same auction scheme is used for the initial allocation of targets to robots. We focus on two design criteria for such coordination systems: First, they have to be computation and communication efficient because they need to coordinate robots in real time. Second, they have to result in a good team performance. We assume in this paper that the team objective is to minimize the sum of the path costs all of robots (for example, the total energy consumed by all robots), where the path cost of a robot is the sum of the edge costs along its path, from its current location to the last target that it visits. We have also studied other team objectives but they cannot guarantee that the team performance is guaranteed to be only a constant factor away from optimum.

Example Coordination Systems

The robots reallocate the unvisited targets among themselves whenever they gain more information about the terrain, each time myopically assuming that their knowledge will not change in the future, that is, that they completely know the terrain. In the remainder of this paper, we study one such cycle of reallocations and thus assume that the terrain is completely known. In this context, we first discuss two common coordination systems with very different properties.

Single-Round Combinatorial Auction

A coordination system based on one single-round combinatorial auction works as follows: Every robot bids on bundles (sets) of targets. It bids the smallest path cost needed to visit all targets in the bundle (from its current location). A central auctioneer determines and informs the winning robots. The winners are determined so as to minimize the sum of the bids of the winning bundles, with the constraint that each robot wins at most one bundle and each target is contained in exactly one bundle. The team performance of a coordination system based on combinatorial auctions is optimal since it takes all positive and negative synergies between targets into account. Two targets are said to exhibit positive (negative) synergy for a robot if their combined cost for the robot is smaller (larger) than the sum of their individual costs. For example, there is a strong positive synergy between two nearby targets because a robot can visit the second target with a much smaller cost after it has reached the first one than from its original location. On the other hand, there is a strong negative synergy between two targets that are on opposite sides of the robot because the robot




Figure 1: Exploration Task 1




Figure 2: Exploration Task 2

can visit the second target only with a much larger cost after it has reached the first one than from its original location. In practice, however, there are three problems when implementing coordination systems based on combinatorial auctions: First, robots cannot bid on all possible bundles of targets because the number of possible bundles is exponential in the number of targets. Second, robots cannot calculate their bids for a given bundle quickly since this requires them to calculate the smallest path cost for visiting a set of targets, which is NP-hard. Third, the winning robots cannot be determined quickly because the winner determination problem is either of exponential size or NP-hard. Researchers have addressed these issues with approximation methods but the resulting coordination systems are complex and no longer able to guarantee optimal team performance.

Parallel Single-Item Auctions

A coordination system based on parallel single-item auctions works as follows: Every robot bids on each target in parallel. It bids the smallest path cost needed to visit the target. The robot that currently owns a target determines and informs the winning robot for the target, which is the robot with the smallest bid on the target. Such a coordination system is simple to implement and computation and communication efficient. Unfortunately, the team performance of a coordination system based on parallel single-item auctions can be highly suboptimal since it does not take any synergies between targets into account. In the example of Figure 1, for example, R1 denotes the location of a robot, and T1, T2 and T3 denote the locations of the targets. (The grid is shown only for the benefit of the reader since we assume that the robots can move in arbitrary directions. However, our observations continue to hold if the robots can move in the 4 compass directions only.) Robot R1 bids 2 on target T1, 2 on target T2, and 3 on target T3 and thus neither models the negative synergy between targets T1 and T2 nor the positive synergy between targets T2 and T3. The following proposition shows that the resulting team performance can be arbitrarily bad.

1. for each robot $r \in R$
2. $T(r) := \emptyset$;
3. while ($T \neq \emptyset$)
4. for each robot $r \in R$
5. for each target $t \in T$
6. $bid(r, t) := PC(r, T(r) \cup \{t\}) - PC(r, T(r))$;
7. submit ($r, t, bid(r, t)$);
8. $(wr, wt) := \arg \min_{r \in R, t \in T} bid(r, t)$;
9. $T := T \setminus \{wt\}$;
10. $T(wr) := T(wr) \cup \{wt\}$;

Figure 3: Sequential Single-Item Auctions (1)

1. for each robot $r \in R$
2. $T(r) := \emptyset$;
3. for each target $t \in T$
4. $bid(r, t) := PC(r, T(r) \cup \{t\}) - PC(r, T(r))$;
5. $bid(r) := \min_{t \in T} bid(r, t)$;
6. $target(r) := \arg \min_{t \in T} bid(r, t)$;
7. submit ($r, target(r), bid(r)$);
8. while ($T \neq \emptyset$)
9. $wr := \arg \min_{r \in R} bid(r)$;
10. $wt := target(wr)$;
11. $T := T \setminus \{wt\}$;
12. $T(wr) := T(r) \cup \{wt\}$;
13. for each target $t \in T$
14. $bid(wr, t) := PC(wr, T(wr) \cup \{t\}) - PC(wr, T(wr))$;
15. for each robot $r \in R$ with $target(r) = wt$
16. $bid(r) := \min_{t \in T} bid(r, t)$;
17. $target(r) := \arg \min_{t \in T} bid(r, t)$;
18. submit ($r, target(r), bid(r)$);

Figure 4: Sequential Single-Item Auctions (2)

Proposition 1 *The sum of all path costs can be an arbitrarily large factor away from optimum when using the coordination system based on parallel single-item auctions in completely known terrain.*

Proof 1 *In the example of Figure 2, the coordination system based on parallel single-item auctions results in $R_i \rightarrow T_i$ (that is, robot R_i visits target T_i) for all $1 \leq i \leq n$ with a sum of all path costs of xn , whereas $R_1 \rightarrow T_1 \rightarrow \dots \rightarrow T_n$ (that is, robot R_1 visits T_1, T_2, \dots, T_n in order) minimizes the sum of all path costs with a sum of all path costs of $n + x - 1$. $\lim_{x \rightarrow \infty} xn / (n + x - 1) = n$ for constant n , and n can be made arbitrarily large. ■*

Sequential Single-Item Auctions

We now discuss the coordination system for which we will show that it combines the advantages of combinatorial auctions and parallel single-item auctions. A coordination system based on a series of single-item auctions, sometimes called sequential (or multi-round) single-item auctions, works as follows: All targets are initially unallocated. Every robot bids on each unallocated target. It bids the increase in its smallest path cost that results from winning the

T4				R1			T1
T3				R2			T2

Figure 5: Exploration Task 3

target that it bids on. The robot with the overall smallest bid is allocated the corresponding target. Each robot simply determines the winning robot quickly for itself by listening to the bids and identifying the smallest bid. Then, each robot re-bids on each unallocated target, and the cycle repeats until all targets are owned by robots. Each robot then calculates the minimum-cost path for visiting all of its targets and moves along this path. Similar ideas of using sequential single-item auctions or ways of calculating the bids have been suggested before in the literature but never been formally analyzed. To describe the coordination system more formally, let R be the set of robots and T be the set of unallocated targets (initially the set of all targets). Let $T(r)$ be the set of targets owned by robot $r \in R$ (initially the empty set). Let $PC(r, T')$ be the smallest path cost for visiting all targets in T' from the current location of robot r . Assume that robot r wins target t . Let $T(r')$ be the set of targets owned by robot $r' \in R$ before robot r wins target t , and $T'(r')$ be the set of targets owned by robot $r' \in R$ afterwards. Thus, $T'(r) = T(r) \cup \{t\}$ and $T'(r') = T(r')$ for all robots $r' \in R$ with $r \neq r'$. Then, the sum of all smallest path costs increases by $\sum_{r \in R} PC(r, T'(r)) - \sum_{r \in R} PC(r, T(r)) = PC(r, T(r) \cup \{t\}) - PC(r, T(r)) = bid(r, t)$. Consequently, a coordination system based on sequential single-item auctions performs hill-climbing since, during each round, a target is allocated to a robot so that the sum of all smallest path costs increases the least. Figure 3 shows a simple implementation where each robot r bids $bid(r, t)$ on each unallocated target t . Figure 4 shows a more complicated implementation where each robot r now bids on at most one of the unallocated targets, namely a target t for which its bid $bid(r, t)$ is smallest. This target is $target(r)$, and robot r bids $bid(r)$ on it. This implementation reduces the number of bids but results in the same allocation of targets to robots as the first implementation (modulo tie breaking) since only the smallest bid of each robot per round can possibly win a target.

A coordination system based on sequential single-item auctions takes some synergies between targets into account but not all of them. In the example of Figure 1, the robot bids 2 on target T1 and 2 on target T2 and thus neither models the negative synergy between targets T1 and T2 nor the positive synergy between targets T2 and T3. However, once robot R1 has won a target, it takes into account the synergy between this target and a target it bids on. For example, assume that robot R1 wins target T2. It then bids 4 on target T1 but only 1 on target T3. Thus, it now models the negative synergy between targets T1 and T2 and the positive synergy between targets T2 and T3. In the example of Figure 5, the coordination system based on combinatorial auctions results

T4				R1			T1
T3				R2			T2

Figure 6: Exploration Task 4

		R1	T1	R2	T2		
--	--	----	----	----	----	--	--

Figure 7: Exploration Task 5

in $R1 \rightarrow T1 \rightarrow T2$ and $R2 \rightarrow T3 \rightarrow T4$ or, if ties are broken differently, $R1 \rightarrow T4 \rightarrow T3$ and $R2 \rightarrow T2 \rightarrow T1$. In both cases, the sum of all path costs is 11. The coordination system based on parallel single-item auctions results in $R1 \rightarrow T1 \rightarrow T4$ and $R2 \rightarrow T2 \rightarrow T3$. The sum of all path costs is 20. The coordination system based on sequential single-item auctions finds the same solution as the coordination system based on combinatorial auctions and thus minimizes the sum of all path costs. However, this is not guaranteed. The exploration task of Figure 6 is similar to the one of Figure 5 but the coordination system based on sequential single-item auctions now finds the same solution as the coordination systems based on parallel single-item auctions and thus does not minimize the sum of all path costs.

Analytical Evaluation

We now evaluate the coordination system based on sequential single-item auctions analytically in completely known terrain under the reasonable assumptions that the costs of moving from location to location are identical for all robots, are symmetrical between locations, and satisfy the triangle inequality.

Lower Bound

One cannot expect coordination systems based on sequential single-item auctions to find optimal solutions of exploration tasks since minimizing the sum of the path costs all of robots is NP-hard for exploration tasks (Lagoudakis *et al.* 2004). In the example of Figure 6, the resulting sum of all path costs is a factor of $20/15 = 1.33$ away from optimum. The following proposition proves a factor that is an even stronger lower bound.

Proposition 2 *The sum of all path costs can be a factor of 1.5 away from optimum when using the coordination system based on sequential single-item auctions in completely known terrain.*

Proof 2 *In the example of Figure 7, the coordination system based on sequential single-item auctions can result in $R2$*

$\rightarrow T1 \rightarrow T2$ with a sum of all path costs of 3, whereas $R1 \rightarrow T1$ and $R2 \rightarrow T2$ minimizes the sum of all path costs with a sum of all path costs of 2. Thus, the resulting sum of all path costs is a factor of $3/2 = 1.5$ away from optimum. (Notice that we rely on a particular tie-breaking behavior in this example but we can easily move the targets by small amounts to guarantee that the robots follow the above paths no matter how ties are broken.) ■

Upper Bound

It might appear that the coordination system based on sequential single-item auctions shares with the coordination system based on combinatorial auctions that the robots cannot calculate their bids quickly since this requires them to calculate the smallest path cost for visiting a set of targets, which is NP-hard. However, the cheapest insertion heuristic can be used to calculate approximate bids quickly, as follows: The robot remembers the path for visiting all targets that it owns already from its previous calculations. It then calculates the path for visiting the target that it bids on plus all targets that it owns already, as follows: It inserts the target that it bids on into all positions on the path for visiting all targets that it owns already, calculates the costs of the resulting paths, and chooses the minimum-cost path. The following theorem is the main analytical result. Bounds for other team objectives can be obtained by similar arguments (Lagoudakis *et al.* 2005).

Theorem 1 *The sum of all path costs is at most a factor of 2 away from optimum when using the coordination system based on sequential single-item auctions in completely known terrain. This result continues to hold if the robots do not calculate the smallest path costs but use the cheapest insertion heuristic to approximate them, resulting in a polynomial-time auction mechanism.*

Proof Sketch 1 *If the sequential single-item auction allocates target t to some robot, then the winning bid is at most twice the cost of the cheapest edge \hat{e} between the set of all previously allocated targets and the set of all yet unallocated targets. This inequality follows from the triangle inequality and cheapest insertion assumption. Consider a hypothetical run of Prim's greedy algorithm, starting with the locations of all robots, that is forced to allocate targets in the same order as the auction by reducing the cost of the cheapest edge between the set of all previously allocated targets and target t to the cost of edge \hat{e} . Combine the optimality property of Prim's algorithm, the inequality above, and monotonicity and triangle inequality properties, including the bounds from the minimum spanning tree to the traveling salesperson problem, to obtain the result. (Lagoudakis *et al.* 2005) contains a formal and complete proof. ■*

The coordination system based on sequential single-item auctions (SSA) combines the advantages of coordination systems based on either combinatorial auctions (CA) or parallel single-item auctions (PSA), as summarized in Figure 8: The implementation of the coordination system based on sequential single-item auctions in Figure 4 runs in polynomial time if the cheapest insertion heuristic is used to calculate

System	Runtime	Number of Bids	Sum of Path Costs
			Minimal Sum of Path Costs
CA	exponential	exponential	1
SSA	polynomial	$ T R $	1.5–2
PSA	polynomial	$ T R $	unbounded

Figure 8: Analytical Worst-Case Results

the bids and results in no more (and likely fewer) bids than a coordination system based on parallel single-item auctions, namely at most $|T| \times |R|$ bids. It is much easier to implement than the coordination system based on combinatorial auctions since the central auctioneer receives exponentially less information and does not need to solve an NP-hard problem; and it provides much better performance guarantees than the coordination system based on parallel single-item auctions. We have also performed experiments that show that the team performance of the coordination system based on sequential single-item auctions is empirically close to optimal (and thus almost as good as what a coordination system based on combinatorial auctions should ideally achieve) and much better than that of the coordination system based on parallel single-item auctions (Tovey *et al.* 2005).

Acknowledgments

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.

References

- Fischer, K.; Muller, J.; Pischel, M.; and Schier, D. 1995. A model for cooperative transportation scheduling. In *Proceedings of the International Conference on Multiagent Systems.*, 109–116.
- 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., and Matarić, M. 2003. Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. In *Proceedings of the International Conference on Robotics and Automation*, 3862–3868.
- Kalra, N.; Stentz, R.; and Ferguson, D. 2004. Hoplitest: A market framework for complex tight coordination in multi-agent teams. Technical Report CMU-RI-TR-04-41, Robotics Institute, Carnegie Mellon University, Pittsburgh (Pennsylvania).
- Lagoudakis, M.; Berhault, M.; Keskinocak, P.; Koenig, S.; and Kleywegt, A. 2004. Near-optimal allocation of exploration tasks to robot teams. In *Proceedings of the International Conference on Intelligent Robots and Systems*.
- Lagoudakis, M.; Markakis, V.; Kempe, D.; Keskinocak, P.; Koenig, S.; Kleywegt, A.; Tovey, C.; Meyerson, A.; and Jain, S. 2005. Auction-based multi-robot routing. In

Proceedings of the International Conference on Robotics: Science and Systems, 343–350.

Nair, R.; Ito, T.; Tambe, M.; and Marsella, S. 2002. Task allocation in the rescue simulation domain: A short note. In Birk, A., and Coradeschi, S., eds., *RoboCup-2001: Robot Soccer World Cup V*, Lecture Notes in Computer Science. Springer.

Simmons, R.; Apfelbaum, D.; Burgard, W.; Fox, D.; Moors, M.; Thrun, S.; and Younes, H. 2000. Coordination for multi-robot exploration and mapping. In *Proceedings of the National Conference on Artificial Intelligence*, 852–858.

Smith, R. 1980. The contract net protocol: high level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29:1104–1113.

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 Telemicroscopy and Telepresence Technologies VII*, volume 4195, 84–95.

Tovey, C.; Lagoudakis, M.; Jain, S.; and Koenig, S. 2005. The generation of bidding rules for auction-based robot coordination. In L. Parker, F. S., and Schultz, A., eds., *Multi-Robot Systems: From Swarms to Intelligent Automata*. Springer. 3–14.

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*, 3016–3023.