

Progress on Agent Coordination with Cooperative Auctions*

Sven Koenig

University of Southern California
Computer Science
skoenig@usc.edu

Pinar Keskinocak and Craig Tovey

Georgia Institute of Technology
Industrial and Systems Engineering
{pinar,ctovey}@isye.gatech.edu

Abstract

Auctions are promising decentralized methods for teams of agents to allocate and re-allocate tasks among themselves in dynamic, partially known and time-constrained domains with positive or negative synergies among tasks. Auction-based coordination systems are easy to understand, simple to implement and broadly applicable. They promise to be efficient both in communication (since agents communicate only essential summary information) and in computation (since agents compute their bids in parallel). Artificial intelligence research has explored auction-based coordination systems since the early work on contract networks (Smith 1980), mostly from an experimental perspective. This overview paper describes our recent progress towards creating a framework for the design and analysis of cooperative auctions for agent coordination.

Introduction

Centralized control is often inefficient for distributed systems in terms of both the required amount of computation and communication since the central controller is the bottleneck of the system. Many researchers have therefore studied agent coordination with cooperative auctions. An auction is “a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants” (McAfee and McMillan 1987). In auction-based coordination systems, the bidders are agents, and the items up for auction are tasks to be executed by the agents. All agents bid their costs. Thus, the agent with the smallest bid cost is best suited for a task. All agents then execute the tasks that they win. Economics has an extensive auction literature but its agents are rational and com-

*We thank Xiaoming Zheng for his overview of auction applications and contributions to the development of algorithms and their experimental evaluation. This material is based on some of our published research on the topic, which was supported by, or in part by, NSF under contract/grant number 0413196, ARL/ARO under contract/grant number W911NF-08-1-0468 and ONR in the form of a MURI under contract/grant number N00014-09-1-1031. 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 or the U.S. government.
Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

petitive, leading to long decision cycles, strategic behavior and possibly collusion. Such issues do not arise in auction-based coordination systems because the agents simply execute their program. Thus, while some insights from economics can be exploited for building auction-based coordination systems (for example, the concept of synergy), many of them do not apply. Conversely, auction-based coordination systems must operate in real-time. Some researchers therefore prefer to use the term “auction-inspired control algorithms” (for decentralized control) over “cooperative auctions” to highlight these differences. This paper provides a unified overview of our progress towards creating a framework for the design and analysis of cooperative auctions for agent coordination, drawing from publications in different venues, including robotics and agents conferences.

Applications

Auction-based coordination systems apply to a wide range of real-time domains.

On-Line Distributed Role Allocation

Allocate roles to agents with different capabilities for the execution of a given plan or playbook so that every task is performed by a qualified agent (Hunsberger and Grosz 2000). Examples include allocating attacker and defender roles to robots in RoboSoccer (Frias-Martinez, Sklar, and Parsons 2004), tasks to ambulance teams, fire brigades and police forces in RoboCup Rescue (Nair et al. 2002), tasks to team members that prevent incursions of aircraft or boats into launch impact zones (Sycara et al. 2005), different observation targets to sensors in a wireless sensor network (for example, for bird habitat sensing) (Yu and Prasanna 2005; Howard and Viguria 2007), and observer and manipulator roles for manipulation tasks (such as in box pushing) (Gerkey and Matarić 2002).

On-Line Distributed Scheduling and Control

Allocate tasks (or processes) to machines (or processors) in a distributed system to minimize latency, maximize throughput or optimize other objectives, possibly considering constraints among the tasks, such as precedence constraints or assignment restrictions. Examples include allocating complex workflows with deadlines to CPUs in grid computing

(Reeves et al. 2005) and on-line control of environmental conditions in smart buildings (Clearwater et al. 1996).

On-Line Distributed Routing

Allocate locations to agents. Examples include allocating accidents to ambulances, incidents to police cars, mines to autonomous underwater vehicles (for de-mining) (Sariel, Balch, and Stack 2006), search-and-rescue locations to first responders, customers to taxi cabs, rocks to Mars rovers (for taking and analyzing rock probes) (Tovey et al. 2005), mines to submarines (for identification) (Sariel, Balch, and Stack 2006), observation locations to robots that map terrain (Simmons et al. 2000) and wings of an art gallery to guard robots (Kalra, Stentz, and Ferguson 2005). In this context, auctions have been used on actual robots (Thayer et al. 2000; Gerkey and Mataric 2002; Zlot et al. 2002).

Testbed: Multi-Robot Routing

The standard testbed of auction-based coordination systems is multi-robot routing (Dias et al. 2006), a special case of on-line distributed routing. For multi-robot routing, the bidders are the robots, and the items up for auction are tasks to visit given targets (locations). The robots are identical and know both their own location and the target locations. They can move and broadcast information without error. They might initially not know where the obstacles are in the terrain but always observe the ones in their vicinity without error. For ease of exposition, we assume in this overview that the robots have to visit all targets (and do not need to return to their initial locations) with a small sum of travel distances. Multi-robot routing problems are similar to traveling salesperson problems and their variants (Lawler et al. 1985), which simplifies their analysis.

Auction-Based Coordination Systems

Minimizing the sum of travel distances is NP-hard for multi-robot routing problems even if the terrain and targets are initially known and do not change (Lagoudakis et al. 2005). Auction-based coordination systems work as follows: Every robot determines a path with the smallest travel distance to visit all of the (unvisited) targets that are allocated to it and starts to move along the path. Thus, the robot does not necessarily visit the targets in the order in which they were allocated to it. Whenever a robot gains more information about the terrain or observes that the terrain has changed, it shares this information with the other robots. If that information increases the travel distance of at least one robot or new targets are introduced, then all robots put their (unvisited) targets up for auction for re-allocation. Each robot then bids in light of the new information, assuming that no changes will be necessary in the future. The auction closes after a predetermined amount of time and the robots are allocated new targets. Auction-based coordination systems based on *parallel auctions* allocate the targets in independent and simultaneous single-round auctions, one for each target. Every robot bids on all targets. They do not take synergies among targets into account, which often results in

a large sum of travel distances.¹ Auction-based coordination systems based on *combinatorial auctions* allocate the targets in one single-round auction. Every robot bids on all bundles (sets) of targets. Hence, all synergies among targets are taken into account, which minimizes the sum of travel distances. However, an exponential (in the number of targets) number of bids must be generated, transmitted, and processed. The approximations necessary to guarantee real-time performance (Berhault et al. 2003) often interact in unpredictable ways and result in complicated code that is difficult to debug, maintain and integrate into robot architectures. Sequential single-item (SSI) auction-based coordination systems have recently emerged as a promising way of combining the advantages of auction-based coordination systems based on parallel auctions (namely, their small number of bids and fast winner determination) and combinatorial auctions (namely, their small sum of travel distances) (Lagoudakis et al. 2005).

Sequential Single-Item (SSI) Auctions

The targets are allocated in one multi-round auction. During each round, every robot bids on each unallocated target, and winner determination then allocates one additional target to one robot. Every robot bids on a target the smallest increase in its travel distance that would result from it being allocated the target that it bids on in addition to all targets allocated to it in previous rounds (marginal-cost bidding) (Sandholm 1993). Winner determination determines the bid with the smallest bid cost and allocates the corresponding target to the corresponding robot (Boutilier, Goldszmidt, and Sabata 1999; Fatima 2006).² SSI auction-based coordination systems take some (but not all) synergies among targets into account and provide the following performance guarantee if the terrain and targets are initially known and do not change. Some intuition for this result can be gained from interpreting the greedy construction of minimum spanning trees as a cooperative auction (Lagoudakis et al. 2004).

Theorem 1 (Lagoudakis et al. 2005) *The sum of travel distances can be a factor of 1.5 larger than minimal but is at most a factor of two larger than minimal, whether each robot calculates its travel distance exactly or uses the cheapest insertion heuristic (Lawler et al. 1985) to determine it approximately, which results in a polynomial-time auction mechanism.*

SSI auction-based coordination systems perform even better experimentally. It can be shown that they perform hill-climbing in each round by allocating one additional target to one robot so that the sum of travel distances increases the least. This insight can be exploited to automatically determine how the robots should bid for different performance

¹Two targets have positive (negative) synergy for a robot if the smallest travel distance for visiting both targets is smaller (larger) than the sum of the smallest travel distances for visiting both targets individually from the current location of the robot.

²In practice, each robot can determine the winning bid quickly itself by listening to the bids and identifying the one with the smallest bid cost.

measures, such as having to visit all targets with a small task-completion or flow time instead of a small sum of travel distances (Tovey et al. 2005). Furthermore, every robot needs to bid on only one target per round, namely the target with the smallest bid cost, since targets with larger bid costs cannot win. The runtime of winner determination during each round is linear in the number of submitted bids, which results in the following desirable communication and runtime complexities of winner determination:

Proposition 1 (Lagoudakis et al. 2005) *The number of bids submitted by all robots during each round and the runtime of winner determination are linear in the number of robots and independent of the number of unallocated targets.*

Improving SSI Auctions

Researchers have investigated several variants of SSI auctions to build SSI auction-based coordination systems that decrease the sums of travel distances while still allocating tasks to robots in real time. Researchers have also investigated how to further improve the target allocation afterwards, for example, using task swaps among robots (Dias and Stentz 2000; Zheng and Koenig 2009).

SSI Auctions with Rollouts

Everything is the same as for SSI auction-based coordination systems, except that every robot now bids on a target the sum of travel distances of the complete target allocation that would result from it being allocated the target that it bids on in addition to all targets allocated to it in previous rounds, all other robots being allocated the targets allocated to them in previous rounds, and then hill-climbing completing this partial target allocation to a complete target allocation (Zheng, Koenig, and Tovey 2006). The bid costs of the robots are now more informed since they are based on complete rather than partial target allocations, which takes more synergies among targets into account and makes hill-climbing less myopic. This is most helpful in the early rounds where the target allocations are far from being complete. Proposition 1 continues to hold trivially. However, the number of rounds increases since each round of the main SSI auction is now preceded by the rounds of the SSI auctions for the corresponding (parallel) rollouts.

SSI Auctions with Bundle Bids

Everything is the same as for SSI auction-based coordination systems, except that every robot now bids on each bundle of at most k unallocated targets, and winner determination then allocates k additional targets to one or more robots, making SSI auctions with bundle bids the same as standard SSI auctions if $k = 1$ and the same as combinatorial auctions if k is large (Koenig et al. 2007). SSI auction-based coordination systems with bundle bids allocate k additional targets to one or more robots in each round so that the sum of travel distances increases the least, which takes more synergies among targets into account and makes hill-climbing less myopic. It can be shown that every robot needs to bid on only a constant number of bundles per round since the other

bundles cannot win. These bundles can be determined automatically. For example, every robot needs to bid on three bundles per round if $k = 2$, namely the single-target bundles with the two lowest bid costs and the double-target bundle with the lowest bid cost, and only seven bundles per round if $k = 3$, which can reduce the number of bids by several orders of magnitude. For example, the number of bids submitted by every robot per round for twenty unallocated targets and $k = 6$ is only 105 instead of 60,459 (Koenig et al. 2007). The winner determination procedure can be automatically determined and, to decrease runtime, be compiled into compact program code (Daniel and Koenig 2009). Proposition 1 continues to hold for fixed k (although the proof is nontrivial), which makes SSI auctions with bundle bids an attractive cross between SSI and combinatorial auctions (Koenig et al. 2007).

SSI Auctions with Regret Clearing

Everything is the same as for SSI auction-based coordination systems, except that winner determination now allocates the target with the largest regret to the robot whose bid cost on it is smallest, where the regret of a target is the difference of the second-smallest and smallest bid cost on it (Koenig et al. 2008). If the terrain and targets are initially known and do not change, SSI auction-based coordination systems with regret clearing provide the following performance guarantee, which could be worse than that of Theorem 1.

Theorem 2 (Koenig et al. 2008) *The sum of travel distances can be a factor of three larger than minimal but is at most a factor of twice the number of targets larger than minimal, whether each robot calculates its travel distance exactly or uses the cheapest insertion heuristic to determine it approximately, which results in a polynomial-time auction mechanism.*

However, SSI auction-based coordination systems with regret clearing can perform much better experimentally than standard SSI auction-based coordination systems, which can be explained as follows: Later allocations of targets to robots are typically more informed than earlier ones since the target allocations in earlier rounds are far from being complete. If a target is allocated to a robot in the current round then one wants to ensure that, if this allocation were postponed, the same allocation would be made in later rounds. SSI auction-based coordination systems with regret clearing achieve this objective by no longer performing hill-climbing in each round but rather maximizing the difference of the second-smallest and smallest team costs that would result from the second-best and best robot, respectively, being allocating one additional target. The communication and runtime complexities of winner determination remain polynomial but Proposition 1 no longer holds since every robot now needs to bid on all unallocated targets per round.

Extensions

Auction-based coordination systems have also been successfully applied to several generalizations of multi-robot routing problems, for example, where a team of robots has to visit a set of given targets with linear decreasing rewards

over time, such as required for the delivery of goods to rescue sites after disasters. The robots have to visit a subset of targets so as to maximize the surplus, which is defined to be the sum of the rewards of the visited targets minus the sum of travel costs. Auction-based coordination systems are able to solve these NP-hard problems in seconds and with a surplus that is comparable to the surplus found by a mixed integer program with a 12 hour time limit (Ekici, Keskinocak, and Koenig 2009). Auction-based coordination systems have also been applied to multi-robot routing problems where a team of robots has to visit a set of given targets with given priorities during given time windows that do not overlap, such as required for planetary exploration. Again, auction-based coordination systems are able to solve these NP-hard problems in seconds and with good team performance (Melvin et al. 2007). Current work includes applying SSI auction-based coordination systems to multi-robot routing problems where robots need to visit some targets simultaneously, such as required for moving large obstacles out of the way cooperatively. Every robot now needs to bid on each pair of time slot and unallocated target. The resulting function maps pairs of time slots and unallocated targets to bid costs but is typically more compact and approximated more easily if it is expressed as a function that maps unallocated targets to reaction functions, where a reaction function maps time slots to bid costs (Zheng and Koenig 2008).

Future Work

There are many alternative approaches for building coordination systems, both centralized (for example, mixed integer programming) and decentralized (for example, distributed constraint optimization or token passing). More work needs to be done on determining when to use which approach since the strengths and weaknesses of the individual approaches are not yet well understood and only a few experimental comparisons exist (Xu et al. 2006). More work needs to be done on developing auction-based coordination systems that better exploit the local (private) information of the agents and auction-based coordination systems for heterogeneous agents. More generally, more work needs to be done on applying auction-based coordination systems in more complex application domains than has been done so far.

- **Example 1:** Consider a heterogeneous team of two different kinds of agents, namely general agents X that can perform both tasks A and B, and specialized agents Y that can perform only task A. Then, it might not be a good idea to let the myopically best agent execute a task (which is what SSI auction-based coordination systems do currently). For example, if two agents X and Y are available and agent X is assigned to execute task A, then no agent is available to execute an arriving task B. On the other hand, if agent Y is assigned to execute task A, then agent X is still available to execute an arriving task A or B. Thus, the agents need to predict the future to achieve a small team cost, perhaps using methods from machine learning (Schneider et al. 2005).
- **Example 2:** Consider tasks that involve time-consuming planning or scheduling to determine the bid costs of the

agents. Then, the agents can calculate only a limited number of bid costs and first need to determine which bid costs to calculate. Thus, they need to predict the bid costs before calculating them, perhaps again using methods from machine learning (Busquets and Simmons 2006).

Finally, more work needs to be done on making auction-based coordination systems robust against error (Sariel, Balch, and Erdogan 2006; Nanjanath and Gini 2008), for example, to ensure that each target gets visited even when robots fail or leave the communication range of other robots.

References

- Berhault, M.; Huang, H.; Keskinocak, P.; Koenig, S.; and Elmaghraby, W. 2003. Robot exploration with combinatorial auctions. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 1957–1962.
- Boutilier, C.; Goldszmidt, M.; and Sabata, B. 1999. Sequential auctions for the allocation of resources with complementarities. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 527–523.
- Busquets, D., and Simmons, R. 2006. Learning when to auction and when to bid. In Gini, M., and Voyles, R., eds., *Distributed Autonomous Robotic Systems 7*. Springer. 21–30.
- Clearwater, S.; Costanza, R.; Dixon, M.; and Schroeder, B. 1996. Saving energy using market-based control. In Clearwater, S., ed., *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific. 253–273.
- Daniel, K., and Koenig, S. 2009. Fast winner determination for agent coordination with SBB auctions [poster abstract]. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 1197–1198.
- 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*, 115–122.
- Dias, M.; Zlot, R.; Kalra, N.; and Stentz, A. 2006. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE* 94(7):1257–1270.
- Ekici, A.; Keskinocak, P.; and Koenig, S. 2009. Multi-robot routing with linear decreasing rewards over time. In *Proceedings of the International Conference on Robotics and Automation*, 958–963.
- Fatima, S. 2006. Sequential versus simultaneous auctions: A case study. In *Proceedings of the International Conference on Electronic Commerce*, 82–91.
- Frias-Martinez, V.; Sklar, E.; and Parsons, S. 2004. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proceedings of the International RoboCup Symposium*, 532–539.
- Gerkey, B., and Mataric, M. 2002. Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation* 18(5):758–768.
- Howard, A., and Viguria, A. 2007. Controlled reconfiguration of robotic mobile sensor networks using distributed

- allocation formalisms. In *Proceedings of the NASA Science Technology Conference*.
- Hunsberger, L., and Grosz, B. 2000. A combinatorial auction for collaborative planning. In *Proceedings of the International Conference on Multi-Agent Systems*, 151–158.
- Kalra, N.; Stentz, R.; and Ferguson, D. 2005. Hoplites: A market framework for planned tight coordination in multi-agent teams. In *Proceedings of the International Conference on Robotics and Automation*, 1170–1177.
- Koenig, S.; Tovey, C.; Zheng, X.; and Sungur, I. 2007. Sequential bundle-bid single-sale auction algorithms for decentralized control. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1359–1365.
- Koenig, S.; Zheng, X.; Tovey, C.; Borie, R.; Kilby, P.; Markakis, V.; and Keskinocak, P. 2008. Agent coordination with regret clearing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 101–107.
- Lagoudakis, M.; Keskinocak, P.; Kleywegt, A.; and Koenig, S. 2004. Auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 1957–1962.
- 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.
- Lawler, E.; Lenstra, J.; Kan, A.; and Shmoys, D., eds. 1985. *The Traveling Salesman Problem*. John Wiley.
- McAfee, P., and McMillan, J. 1987. Auctions and bidding. *Journal of Economic Literature* 15:699–738.
- Melvin, J.; Keskinocak, P.; Koenig, S.; Tovey, C.; and Ozkaya, B. Y. 2007. Multi-robot routing with rewards and disjoint time windows. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 958–963.
- 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.
- Nanjanath, M., and Gini, M. 2008. Repeated auctions for robust task execution by a robot team. Technical Report 08-032, University of Minnesota, Computer Science Department, Minneapolis (Minnesota).
- Reeves, D.; Wellman, M.; MacKie-Mason, J.; and Osepashvili, A. 2005. Exploring bidding strategies for market-based scheduling. *Decision Support Systems* 39:67–85.
- Sandholm, T. 1993. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the International Workshop on Distributed Artificial Intelligence*, 295–308.
- Sariel, S.; Balch, T.; and Erdogan, N. 2006. Robust multi-robot cooperation through dynamic task allocation and precaution routines. In *Proceedings of the International Conference on Informatics in Control, Automation and Robotics*, 196–201.
- Sariel, S.; Balch, T.; and Stack, J. 2006. Empirical evaluation of auction-based coordination of AUVs in a realistic simulated mine countermeasure task. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems*, 197–206.
- Schneider, J.; Apfelbaum, D.; Bagnell, D.; and Simmons, R. 2005. Learning opportunity costs in multi-robot market based planners. In *Proceedings of the International Conference on Robotics and Automation*, 1151–1156.
- 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 AAAI 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.
- Sycara, K.; Scerri, P.; Srinivas, S.; and Lewis, M. 2005. Task allocation in teams for launch and range operations. In *Proceedings of the International Conference on Human Computer Interaction*.
- 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, 84–95.
- Tovey, C.; Lagoudakis, M.; Jain, S.; and Koenig, S. 2005. The generation of bidding rules for auction-based robot coordination. In Parker, L.; Schneider, F.; and Schultz, A., eds., *Multi-Robot Systems: From Swarms to Intelligent Automata*. Springer. 3–14.
- Xu, Y.; Scerri, P.; Sycara, K.; and Lewis, M. 2006. Comparing market and token-based coordination [short paper]. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 1113–1115.
- Yu, Y., and Prasanna, V. 2005. Energy balanced task allocation for collaborative processing in wireless sensor networks. *Mobile Networks and Applications* 10(1-2):115–131.
- Zheng, X., and Koenig, S. 2008. Reaction functions for task allocation to cooperative agents. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 559–566.
- Zheng, X., and Koenig, S. 2009. K-swaps: Cooperative negotiation for solving task-allocation problems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 373–379.
- Zheng, X.; Koenig, S.; and Tovey, C. 2006. Improving sequential single-item auctions. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 2238–2244.
- 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.