

Market-Based Algorithms for Allocating Complex Tasks*

Xiaoming Zheng and Sven Koenig

Department of Computer Science
University of Southern California
Los Angeles, CA 90089-0781
{xiaominz,skoenig}@usc.edu

Introduction

It is often important to coordinate teams of cooperative agents in a distributed manner. We study how to assign tasks to cooperative agents so that the resulting team cost is small (that is, team performance is high). Market-based mechanisms are promising distributed task-allocation methods. Robotics researchers have recently studied how to use sequential single-item (SSI) auctions to allocate simple tasks to agents (Tovey et al. 2005), where simple tasks need to be performed by exactly one agent. In this paper, we study how to allocate complex tasks to agents, where complex tasks (different from (Zlot 2006)) need to be performed by several agents simultaneously. Our **motivating example** is *multi-agent routing*, where the tasks are to visit targets in the plane. Each target needs to be visited by a given number of agents that depends on the target. Simple targets need to be visited by only one agent, while complex targets need to be visited by several agents simultaneously. For example, large fires can be extinguished only with several fire engines, and heavy objects can be moved only with several robots. The agent cost is the smallest sum of travel and wait times needed by an agent to visit all targets assigned to it. The team cost is the largest agent cost of any agent (that is, the task-completion time). Our objective is to determine which targets each agent should visit and when it should visit them so that the team cost is minimal. Most **existing research** treats task allocation as a complete set partitioning or set covering problem where the cost of performing a complex task with a given set of agents is either pre-defined (Abdallah and Lesser 2004) or easy to calculate (Shehory and Kraus 1998). In the context of multi-agent routing, however, the agent cost of visiting an additional complex target depends not only on the target and the agent that visits it but also on the visit time. The most closely related research to our work is some of our own previous research for the special case where each agent can visit at most one complex target (Zheng and Koenig 2008).

*This material is based upon research 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 form of a MURI under contract/grant number N00014-09-1-1031. Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Our Approach

We intend to develop auction-like algorithms for the allocation of complex tasks, similar to SSI auctions for the allocation of simple tasks. *SSI auctions* assign simple tasks to agents in multiple rounds. In each round, each agent bids on each unassigned task the minimal increase in its agent cost in case it has to perform this task in addition to all tasks already assigned to it in previous rounds. The auctioneer then assigns exactly one unassigned task to one agent, namely the task to the agent corresponding to the bid with the smallest bid cost, which performs hill-climbing and thus increases the team cost the least. In the context of multi-agent routing, however, SSI auctions cannot be used directly for the allocation of complex targets because i) agents cannot just submit one numerical bid per complex target since their bid cost depends on the visit time of the complex target; and ii) the auctioneer cannot just determine the bid with the smallest bid cost since it has to solve a non-trivial scheduling problem to determine the visit times of the complex targets so that the team cost increases the least. In this abstract, we suggest solutions to both of these problems.

Characterizing Agent Costs

Assume for now that only **one complex target** x is assigned to agent a , possibly in addition to several simple targets (Zheng and Koenig 2008). Since the complex target needs to be visited by several agents simultaneously, agent a cannot determine the visit time of target x by itself. It can communicate its local knowledge with respect to target x to the auctioneer by stating its agent costs for visiting target x for all visit times t , written as $\mathcal{F}_a^x(t)$. We now discuss how to calculate this *reaction function* $\mathcal{F}_a^x(t)$. Since the agent cost is the smallest sum of travel and wait times needed by the agent to visit all targets assigned to it, we first consider the travel time of the agent as follows: For each possible visit time t of complex target x , we find the smallest travel time of agent a in case it has to visit complex target x at visit time t and all simple targets assigned to it at any visit times of its choosing but is not allowed to wait at any target. The travel time is infinity if this restriction cannot be satisfied. Let P_a denote the set of visit times of the complex target whose corresponding travel times are finite. For each visit time $t_p \in P_a$, let c_p denote the corresponding travel time and o_p denote the corresponding visit order of all (simple and complex) targets

assigned to the agent.¹ We now include the wait time of the agent as follows: Assume that agent a is required to visit complex target x at visit time t . Without loss of generality, we assume that agents wait only at complex targets and leave them immediately after their visit. For each visit time $t_p \in P_a$ with $t_p \leq t$, visiting the targets in visit order o_p is one way how agent a can visit complex target x at visit time t . The resulting sum of travel and wait times is $c_p + t - t_p$, where c_p is the travel time and $t - t_p$ is the wait time at the complex target. The agent cost is the smallest sum of travel and wait times. Thus, $\mathcal{F}_a^x(t) = \min_{t_p \in P_a, t_p \leq t} (c_p + t - t_p)$. The key idea here is to express a reaction function whose domain is infinite with a finite number of values. We now generalize this idea to the case where **multiple complex targets** are assigned to agent a , possibly in addition to several simple targets. Let $\mathbf{x}_a = (x_a^1, \dots, x_a^k)$ denote a vector of complex targets assigned to agent a , where the complex targets are visited by the agent in the order specified by \mathbf{x}_a . For each possible vector $\mathbf{t}_a = (t_a^1, \dots, t_a^k)$ of visit times of the complex targets in \mathbf{x}_a with $t_a^1 \leq \dots \leq t_a^k$, we find the smallest travel time of agent a in case it has to visit the complex targets \mathbf{x}_a at visit times \mathbf{t}_a and all simple targets assigned to it at any times of its choosing but is not allowed to wait at any target. The travel time is infinity if this restriction cannot be satisfied. Let P_a denote the set of vectors of visit times of the complex targets whose corresponding travel times are finite. For each vector $\mathbf{t}_p \in P_a$ of visit times, let c_p denote the corresponding travel time and o_p denote the corresponding visit order of all (simple and complex) targets assigned to the agent. Assume that agent a is required to visit the complex targets \mathbf{x}_a at visit times \mathbf{t}_a . For each vector $\mathbf{t}_p \in P_a$ of visit times with $0 \leq t_a^1 - t_p^1 \leq \dots \leq t_a^k - t_p^k$, visiting the targets in visit order o_p is one way how agent a can visit the complex targets \mathbf{x}_a at visit times \mathbf{t}_a . The resulting sum of travel and wait times is $c_p + t_a^k - t_p^k$, where c_p is the travel time and $t_a^k - t_p^k$ is the total wait time at all complex targets. The agent cost is the smallest sum of travel and wait times. Thus, $\mathcal{F}_a^{\mathbf{x}_a}(\mathbf{t}_a) = \min_{\mathbf{t}_p \in P_a, 0 \leq t_a^1 - t_p^1 \leq \dots \leq t_a^k - t_p^k} (c_p + t_a^k - t_p^k)$ for the visit order of the complex targets specified by vector \mathbf{x}_a . The reaction functions for all possible visit orders of the complex targets assigned to an agent then specify its agent costs for visiting its complex targets for all visit times.

Determining Optimal Visit Times

Assume again that only **one complex target** is assigned to each agent, possibly in addition to several simple targets (Zheng and Koenig 2008). Assume further that the auctioneer has to determine when a set of agents A should visit the complex target x assigned to all of them so that the team cost is minimal. Per definition, the optimal visit time is $\min_{t \in \mathbb{R}_+} \max_{a \in A} \mathcal{F}_a^x(t)$. It can be determined in finite time as follows: Let $T = \cup_{a \in A} P_a$ denote the set of visit times

¹ P_a can be constructed in finite time as follows: First, construct all possible visit orders of all (simple and complex) targets assigned to the agent. Second, for each visit order, find the visit time of the complex target in case the agent visits all targets in the given visit order without waiting at any target. P_a is the set of all such visit times. Third, if more than one visit order results in the same visit time, keep only the one with the smallest travel time.

of the complex target such that at least one agent does not wait at the complex target. Then, the optimal visit time is $\min_{t \in T} \max_{a \in A} \mathcal{F}_a^x(t)$. This can be seen as follows: For any visit time $t \notin T$ of the complex target, all agents have to wait at complex target x . Let s denote the smallest wait time of any agent. Then, the team cost decreases by s if all agents visit complex target x at visit time $t' = t - s$, where $t' \in T$. Thus, the previous team cost was not minimal. The key idea here is to calculate the minimization over an infinite number of potential visit times by considering only a finite number of values. We now generalize this idea to the case where **multiple complex targets** can be assigned to each agent, possibly in addition to several simple targets. Assume that the auctioneer has to determine when a set of agents A should visit the complex targets assigned to them so that the team cost is minimal. Let $\mathbf{x} = (x^1, \dots, x^n)$ denote the vector of complex targets assigned to the set of agents, where each complex target is assigned to at least one agent and the complex targets are visited in the order specified by \mathbf{x} . Let T denote the set of vectors $\mathbf{t} = (t^1, \dots, t^n)$ of visit times of the complex targets such that, for each complex target, at least one agent does not wait at the complex target. Then, $\min_{\mathbf{t} \in \mathbb{R}_+^n} \max_{a \in A} \mathcal{F}_a^{\mathbf{x}_a}(\mathbf{t}_a) = \min_{\mathbf{t} \in T} \max_{a \in A} \mathcal{F}_a^{\mathbf{x}_a}(\mathbf{t}_a)$, where \mathbf{x}_a and \mathbf{t}_a are the projections of \mathbf{x} and \mathbf{t} onto the complex targets assigned to agent a and their visit times, respectively.

Conclusions

We extended previous research to the allocation of complex tasks, where complex tasks need to be performed by several agents simultaneously. We proposed reaction functions to succinctly characterize the agent costs for executing complex tasks at given times and illustrated an important property of the execution times of complex tasks that minimize the team cost, which facilitates winner determination by the auctioneer. It is future work to build an agent-based coordination system similar to SSI auctions that exploits this property.

References

- Abdallah, S., and Lesser, V. 2004. Organization-based cooperative coalition formation. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, 162–168.
- Shehory, O., and Kraus, S. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence* 101(1–2):165–200.
- 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.
- Zheng, X., and Koenig, S. 2008. Reaction functions for task allocation to cooperative agents. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 559–566.
- Zlot, R. 2006. *An Auction-Based Approach to Complex Task Allocation for Multirobot Teams*. Ph.D. Dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.