

TIERED AUCTIONS FOR MULTI-AGENT COORDINATION IN DOMAINS WITH PRECEDENCE CONSTRAINTS

E. Gil Jones*, M. Bernardine Dias, and Anthony Stentz
Robotics Institute, Carnegie Mellon University
Pittsburgh, PA, 15213

ABSTRACT

Many applications require teams of robots to cooperatively execute complex tasks. Among these domains are some that require robots to interact closely at particular times and locations to accomplish some task components, but otherwise allow the team members to act independently. Successful execution in such domains often requires agent interactions that must adhere to constraints of precedence. Precedence requirements can occur when agents' plans call for certain pre-conditions to be met at particular times and places. In this work we focus on precedence-constrained emergency response. In this domain a group of fire trucks agents attempt to navigate through a city in order to extinguish a set of fires that have occurred in the wake of a large-scale disaster. Another effect of the disaster is that debris have blocked roads in the city, making roads impassable for the fire trucks. Debris can be cleared by bulldozer robots, which are also operating in the environment. To maximize fire fighting performance fire trucks and bulldozers must determine when and where debris clearance interactions should occur and who should be involved.

Our proposed method for coordination in domains with precedence-constrained interactions is a market-based approach to planning, allocating, and scheduling that uses a novel tiered auction framework. The tiered auction framework allows agents to solicit the assistance of other agents in determining their suitability for a task; in this framework agents hold sub-auctions to decide what interactions may best address application constraints. For the emergency response domain we propose a tiered auction method that uses single-task fire truck assignment at the top tier and multi-task bulldozer assignment at the second auction tier so that fire trucks can find maximally efficient routes to fires. We show that the tiered auction approach improves over a standard single-tier market-based approach in a simulated emergency-response domain.

1. INTRODUCTION

Research efforts in robotics have increasingly turned towards using teams of robots to collectively address tasks

rather than employing a single robot operating independently. A team of robots acting together can often outperform single robot solutions in terms of quality and robustness; at the same time, equipping multiple robots to collectively address application tasks requires addressing a substantial research challenge of coordinating the efforts of the robots. To mitigate the difficulties associated with coordinating a team of robots research in multi-robot systems has centered around approaches where agents each work by themselves. These approaches are best suited to domains with independent tasks, where the efforts of single robots acting independently are sufficient. This work concerns itself, however, with domains in which agents must work together more closely; we are particularly interested in domains that require that some agents satisfy precedence constraints to enable other agents to address domain tasks. Coordinating agents in domains with precedence constraints involves addressing substantial new challenges associated with the interdependence of agents' schedules.

In this work we focus on a single motivating domain: precedence-constrained emergency response. In this domain a team of robots capable of extinguishing fires is operating in a city that has been ravaged by a disaster of significant proportions. These fire truck robots are tasked with moving to various locations around the city to extinguish fires that have been reported by an autonomous air vehicle; new fires are frequently discovered. Time is of the essence, and the robots must move to the locations of the fires and extinguish them as quickly as possible. Unfortunately, the disaster has rendered many of the roads in the city impassable; they are covered with debris and wreckage, creating obstacles around which the bulky fire extinguishing robots cannot navigate. Suppose that there is another group of bulldozer robots in the team that are designed to move freely in rough terrain, and can clear roads of wreckage and debris. The bulldozer agents can assist the fire truck agents by satisfying debris clearing preconditions along routes to fires. We suppose that there is a global objective function that is a sum over time-decreasing rewards offered for extinguishing fires. Addressing the coordination problem posed by such a domain so as to maximize global objective reward requires not only assigning fire trucks to fires, but also selecting routes that the trucks

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE Tiered Auctions for Multi-Agent Coordination in Domains with Precedence Constraints				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University, Robotics Institute, Pittsburgh, PA, 15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

will take through the city and assigning bulldozers to clear debris along those routes.

In this work we propose a novel method, tiered auctions, that enables us to efficiently search the space of inter-dependent agent plans to determine allocations and schedules that maximize an objective function associated with the emergency response domain. Our method expands on existing auction-based coordination methods for domains with independent tasks by equipping agents with the ability to solicit assistance during allocation to account for task constraints. We will show that our method outperforms a method that allocates tasks without considering precedence constraints in a simulated precedence constrained emergency response domain.

We next address related work, and then describe our approach in greater detail. Experimental results and our conclusions follow.

2. RELATED WORK

Task allocation for tasks that are not inter-related by constraints has been widely studied. TraderBots uses a market-based approach with sequencing for a wide variety of domains with different independent tasks and objective functions [Dias et al., 2004] [Zlot et al., 2002]. Berhault et al. use a combinatorial auction approach for time-extended assignment of exploration tasks with a travel cost metric [Berhault et al., 2003]. Zlot [Zlot, 2006] focuses on domains with complex tasks that must be broken down into independent tasks for execution. As we will show, task allocation methods in precedence-constrained domains that do not consider constraints between tasks do not perform well in comparison to our proposed method.

Some previous work has explicitly consider domains where tasks are inter-related by constraints [MacKenzie, 2003],[Lemaire et al., 2004]. In neither work do they propose a mechanism for searching a large space of possible plans for tasks; in our domains of interest fire trucks need to efficiently search a huge space of possible route plans.

The final relevant area of work is associated with multi-robot coordination for domains where tasks require the joint efforts of multiple agent. Approaches to these domains are generally called coalition formation algorithms, as they seek to form coalitions of agents to address tasks [Vig and Adams, 2006] [Tang and Parker, 2007] [Jones et al., 2006] [Sariel et al., 2007]. Koes et al. use an MILP-based formulation for a search and rescue domain where tasks have associated rewards that decay linearly, and the system can be constrained in a variety of ways including task ordering, time-oriented task constraints, capability constraints, and global resource constraints [Koes et al., 2005]. While all of these methods contain

mechanisms for coordinating given the presence of inter-task constraints, all are computationally intensive and not well-suited for efficiently searching the large space of possible fire truck routes and associated debris requirements associated with our domains of interest. We leave comparison of our methods with these methods in terms of performance and computation as future work.

3. METHODS

In this section we detail our approach to multi-agent coordination for domains with precedence constraints. One reasonable approach we could take is to use a minimally modified version of a method from our previous work, where we used auctions for coordination in a fire-fighting emergency response domain with deadlines but no precedence constraints (hence no need for bulldozers); our sequential single-item auction approach is described in [Jones et al., 2007]. In our approach when emergency tasks are discovered they are passed to a central dispatcher/auctioneer (D/A), who keeps a list of tasks to be allocated. This agent periodically announces a task auction and sequentially auctions tasks to all available emergency response agents. Each task is assigned to the fire truck agent that determines that it can generate the most additional value in terms of an objective score for performing a task given that it already may have a number of task assignments. One possible way to adapt this approach to our precedence-constrained emergency response domain would be to have fire trucks bid for tasks based on calculating shortest path distances to new fire tasks; this approach would not have fire trucks reason about debris along paths or bulldozer schedules and allocations. Once fire tasks had been assigned the fire trucks would then use some method to recruit bulldozers to debris tasks as needed. We call this the “allocate-then-coordinate” approach. While this approach is straight-forward and would require minimal modifications to our previous approach, we think it will achieve low performance in scenarios with more than a few debris, as debris density and bulldozer availability become increasingly important factors in solution quality.

Our proposed method for allocation in precedence-constrained domains is a tiered auction approach where agents try to reason simultaneously about allocation and route planning. In order to solicit the contributions of other agents during the bidding process we must extend the single-tiered auction approach: our system for emergency response coordination uses two auction tiers. At the top tier fire tasks are auctioned to fire trucks using single-item sequential auctions. In our approach trucks only bid on fires if they are currently idle, having completed their last assigned fire. We use this approach – called instantaneous allocation in accordance with the widely-used multi-agent approach nomenclature proposed by Gerkey and Mataric [Gerkey and Mataric, 2004] – in order to limit problem

complexity; multiple fire assignment is left as future work. In order to bid on fires trucks must search through the space of possible routes to the fire; to assess a route fire truck can hold sub-auctions to recruit bulldozer assistance to clear debris along the route. These sub-auctions comprise the second tier of auctions. Fire trucks may need to hold a series of sub-auctions for different routes to determine the route which will allow them to reach the fire being auctioned as quickly as possible.

While assigning a single fire per fire truck can substantially reduce problem complexity, assigning only a single debris pile per bulldozer could potentially result in poor performance, especially for domains with dense debris, as there may be more debris along desirable routes than there are available bulldozers. Thus while we use instantaneous assignment for fire truck allocation, our approach supports assigning multiple tasks to bulldozers - we use a partially time-extended approach in the Gerkey and Mataric parlance. In our approach bulldozers can be assigned an arbitrary number of tasks, but will only bid on tasks in debris sub-auctions based on adding tasks to the end of their schedules. As we do not equip bulldozers to alter the order of assigned tasks in their schedules or insert new tasks at arbitrary points in their schedules, our approach is only partially time-extended. Fully time-extended bulldozer allocation is left as future work.

We first describe our approach at a high level, and then illustrate approach details using an example scenario. We conclude our methods session with a discussion of bounding approaches we use to limit the search space of possible coordinated schedules.

3.1 Approach Overview

The goal of the auction process is to assign each idle fire truck a single fire. During the assignment process fire trucks determine which route to take to the fire and assign each debris pile along the route to a bulldozer. The auction process proceeds as follows:

1. The dispatcher/auctioneer (D/A) generates a list of unallocated fires that can potentially be assigned to idle agents.
 - (a) The D/A announces an auction for the task at the top of the list. An auction call for that fire F is sent to each fire truck agent.
 - (b) All idle fire trucks enter into an sub-auction loop:
 - i. An A* algorithm returns a path P from a truck's current position to the fire being auctioned.

- ii. Each truck considers each debris pile along P in the order that they will be encountered when following P .
- iii. Some debris may already be allocated to bulldozers as part of some existing commitment. The truck confers with the D/A to get information on the scheduled completion time of these debris piles.
- iv. If some debris have not been previously allocated, the fire truck generates a new auction context C associated with P - this context is associated not only with the fire F but with P , a particular route for the fire truck to reach F .
- v. For each unallocated debris D the fire truck runs a sub-auction:
 - A. A call for bids for D is sent to each bulldozer. The call contains the auction context, the debris ID, and the debris location.
 - B. If a bulldozer does not have a previous schedule associated with C it creates a new schedule associated with C - this schedule is just a copy of its current schedule.
 - C. Each bulldozer adds D to the end of its schedule associated with C , creating a candidate schedule.
 - D. The bulldozers then bid their completion time for D based on their candidate schedules.
 - E. The fire truck awards the auction to the bid corresponding to the earliest task completion time.
 - F. The winning bulldozer adopts the candidate schedule as its new context schedule for C .
 - G. The losing bulldozers discard their candidate schedules, and will bid on future tasks using their previous C context schedules.
 - H. Each unallocated debris along P is allocated in this manner.
- vi. When all unallocated debris have been allocated, trucks can compute their completion time for F . If the reward for F is better than the best previously considered route it becomes the leading candidate schedule. Otherwise the schedule and all associated bulldozer schedules can be discarded.
 - (c) The sub-auction loop concludes when no paths to F could offer a faster completion time.
 - (d) The fire trucks then place bids for F based on scheduled task reward.

2. A task is awarded to the single fire truck that can achieve the highest reward given all bids for all tasks.
3. The winning fire truck adopts the schedule it has stored for the task it has won and informs bulldozers to adopt their context schedules associated with the winning bid.
4. If any fire trucks remain untasked, the auction loop continues.

After the auction process concludes all agents can then execute their schedules. The auction process repeats when any fire truck becomes idle.

3.2 Example scenario

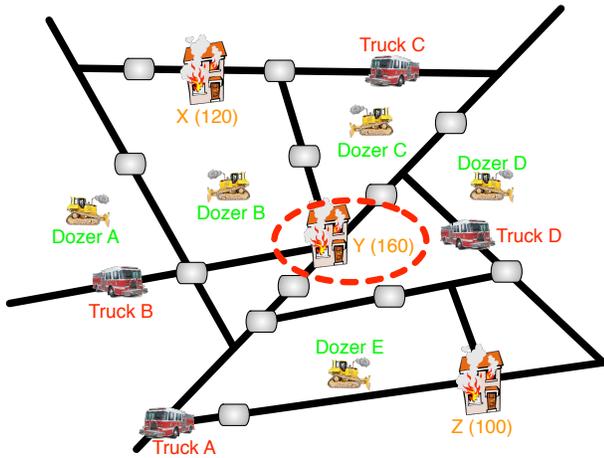


Figure 1: Example scenario in the emergency response domain with four fire trucks, four bulldozers, and three fires. Current fire rewards are given in parentheses. Black lines represent roads, and debris are represented by oval boxes. Fire Y is the first fire selected for auction.

We will now walk through an example of our auction procedure in order to explicate the function of our system. In Figure 1 we show a scenario with 4 fire trucks and 4 bulldozers. We suppose that the D/A has generated a list of available fires and that fire Y is at the top of the list and will be offered for auction first. We assume that all trucks are idle and will bid on the fire.

We will focus on the planning procedure for truck A. We suppose that the truck's A* algorithm returns the shortest distance path for consideration first. This path is shown in blue in Figure 2. Truck A determines that there are unallocated debris along this path and generates a new context 10001 associated with this particular route for reaching Y. It then uses a sub-auction to solicit bulldozer assistance for the first debris task in the route. We assume that all bulldozers are idle; in this case they generate empty schedules for

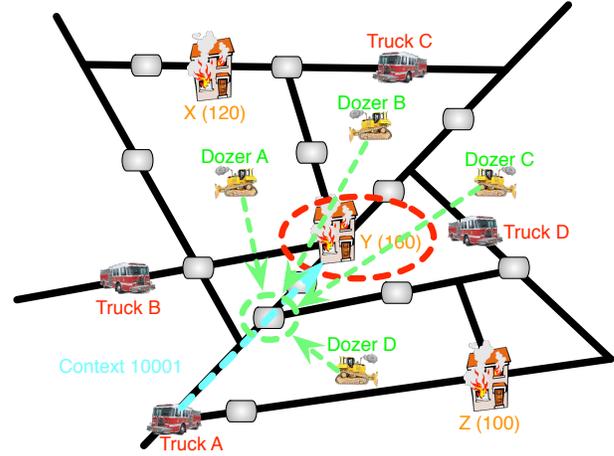


Figure 2: Fire truck A holds a sub-auction for its shortest path route (the blue dotted line) to fire Y, which passes through two unallocated debris. Truck A creates context 10001 for the route to Y, and holds a sub-auction for the first debris pile along the route. Bulldozers bid based on shortest path routes to the debris.

the new context and add the task to their empty schedules, bidding their completion times for the task. Bulldozer D wins the sub-auction, and replaces its empty context schedule for 10001 with the new context schedule with the debris task - all other bulldozers discard their candidate schedules and maintain their empty context schedules. The fire truck then sub-auctions the next task along its path using the same 10001 context designation. Now each bulldozer bids based on adding the task to the end of their context schedules; bulldozer D bids for the new debris task based not on having an empty schedule but instead based on having already won a task for this context. The resulting bulldozer schedules are shown in Figure 3. The sub-auction is awarded to bulldozer A. Note that bulldozer D was actually the closet bulldozer to both debris piles, but that assigning both debris tasks to D would cause more delay to truck A than the allocation that resulted using contexts. With all debris for this particular route allocated truck A can actually compute its completion time and associated reward for fire Y given that this route would be taken. The computation of the completion time is shown in Figure 4. Now truck A must consider other routes to the fire; for each route under consideration it generates a new context and holds a sub-auction for all unallocated debris. Figure 5 shows the results of following another route to Y - taking this route results in a slightly higher reward for Y based on a 2 cycle improvement in completion time. Note that while this route is longer in distance it has fewer debris and in this instance results in a faster completion time for the fire.

Each agent places a bid for fire Y and reports the results to the D/A, who then places a call for bids for fire X, the next fire in the D/A's list. When all bids have been

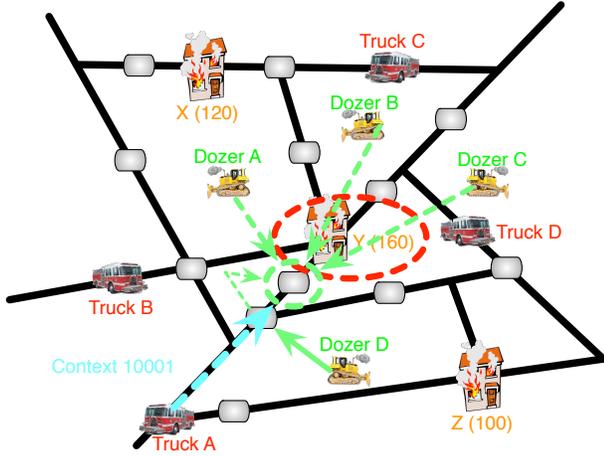


Figure 3: Bulldozer schedules used for bidding for the second debris in truck A’s route to fire Y. Note that Bulldozer D’s schedule for the second debris reflects that it has already won a previous debris sub-auction.

received for all fires, the D/A awards the single best bid among all bids received for all fires. The fire truck with the winning bid then adopts the associated context and informs the bulldozers involved in the plan that they should adopt their context schedules as well. Bulldozers that have previously been assigned debris will then only bid for new debris tasks based on schedules that reflect their commitments to clear the previously allocated debris. The process then repeats until all fires have been allocated or all fire trucks have been assigned a task. The resulting allocation for this scenario is shown in Figure 6.

3.3 Bounding for faster computation

We have implemented a number of refinements in the above algorithm with the goal of making the auctioning process faster and more efficient without sacrificing solution quality. Our bounding methods stem from the observation that during an auction round only a single task will be awarded. If at any time, then, an agent determines that a fire or route to a fire could not possibly generate higher reward than one that it has previously considered during the same auction round it can stop evaluating that fire or route. We incorporate this observation into bounding into all stages of the route evaluation process, including path planning and the sub-auction process. Thus in practice agents generally search only a small fraction of the full space of all possible paths to all auctioned fires. By employing frequent bounding and heuristics we can search the space more quickly without sacrificing solution quality.

4. RESULTS

Before we explore results comparing our tiered auction allocation approach with an “allocate-then-

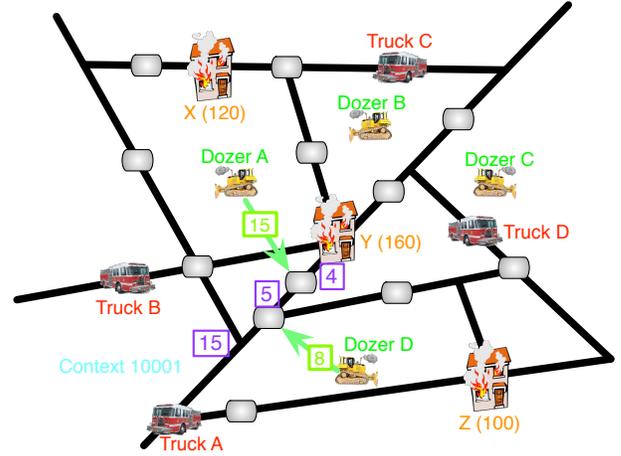


Figure 4: With all debris allocated the fire truck can compute its arrival time and resulting award for the fire Y. Purple boxes shown path lengths from truck A’s current location to the first debris pile, from the first debris pile to the second debris pile, and from the second debris pile to the fire Y. The bulldozer D takes 8 cycles to reach the debris, which takes 6 cycles to clear, and so will be cleared at time 14. Thus once the fire truck reaches the first debris pile at time 15 it has already been cleared, and it can proceed directly to the second debris pile, which it reaches at time 20. Bulldozer A takes 15 cycles to reach the second debris pile, and 8 cycles to clear it. Thus the fire truck must wait until time 23 before proceeding past the second debris pile. It will leave the second debris pile at time 23 and reach the fire at time 27. The fire takes 5 cycles to extinguish, meaning that taking this route will result in truck A extinguishing the fire after 32 cycles. The fire decays at 1.5 units/cycle from an original reward of 160, thus yielding a scheduled reward of 112 ($160 - (32 * 1.5)$).

coordinate”(ATC) approach we must describe our simulated experimental setup. Our fire truck agents operate in a grid network of streets - truck planning is doing using a graph-based map representation. Bulldozers plan using a grid-based map representation. In our simulated domain we suppose that fire discovery is the product of a single Poisson process, the standard distribution used in queuing theory to represent stochastic arrival times of independent tasks. The parameter λ for the Poisson process represents the expected rate of task issuance. Our program is currently set up to directly compare two different approaches; to make the comparison as accurate as possible we have agents operating under the two different approaches operate on exactly the same domain instance: agents begin in the same locations, randomly-generated debris are placed in identical locations, and the same set of fires are discovered at the same time between the two approaches. The only differences between the testing environments are caused by the differences in coordination methods.

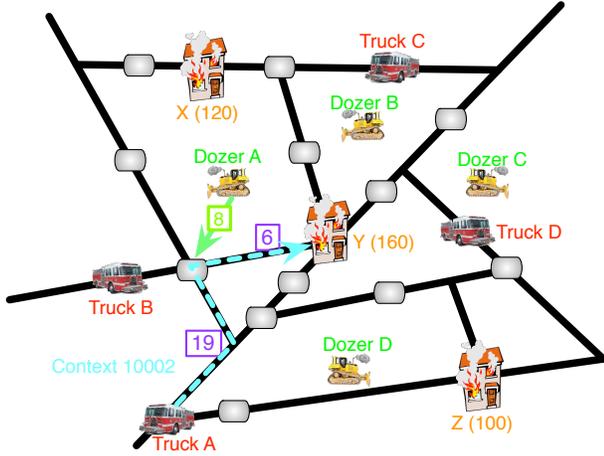


Figure 5: A different context yields a better route to fire Y. Path lengths for the fire truck shown in purple boxes. Dozer A clears the debris at time 15, allowing the truck to reach Y at 25 cycles and extinguish it at 30 cycles, yielding a reward of 115.

There are a number of parameters that go into forming a particular instance of the domain. In our experiments we keep most of the parameters constant: we run our experiments with three fire trucks and three bulldozers operating in a 4-by-6 city network. The composition of the buildings in the city blocks, which the bulldozers must avoid when moving to debris locations, are generated randomly for each domain instance. In each domain instance there were 10 known fires at time 0, and the λ value associated with the Poisson process was set to 1, meaning an expected issuance rate of 1 fire per time cycle. Once at a fire site a fire truck will be able to extinguish the fire in a single time unit independent of fire size. We assume that fires have an initial value that is drawn from a Gaussian distribution with a mean of 3000 and a standard deviation of 100, and that the reward for a fire linearly degrades at a rate of 25 units/cycle for all fires.

We now present the results of a comparison of our tiered auction approach, which considers precedence interactions during allocation and an ATC approach which first allocates without considering precedence interactions and then determines coordination after allocation. We implemented the ATC approach using some of the components of our tiered auction system. We use a single auction tier for fire allocation; fire trucks bid on fires based on shortest path schedules to fires without holding sub-auctions or considering debris along the routes in any way. Once a truck has been awarded a fire it then searches among possible paths to reach the fire using the same procedure used in our tiered auction, holding auctions among the bulldozers as necessary. This approach essentially uses two separate single-tier auctions instead of our multi-tiered auction system.

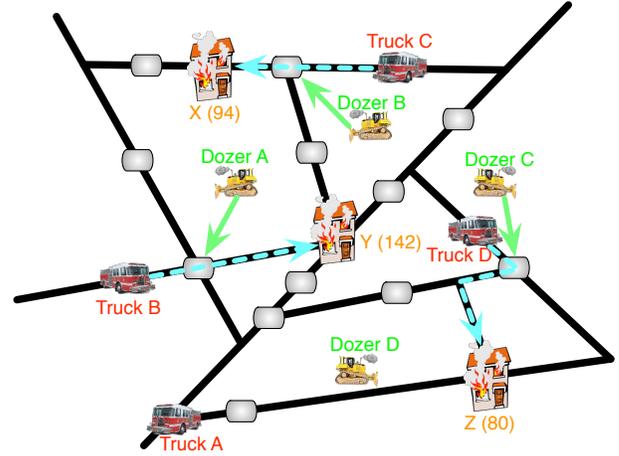


Figure 6: All fires have been assigned. Truck paths are shown with blue dotted lines and green solid lines show bulldozer paths to their assigned debris. Fire rewards at their scheduled completion times are shown in parentheses.

We expect the performance difference between the ATC approach and our tiered auction approach to be most pronounced in domain instances where precedence interactions are required frequently. As negotiating precedence interactions becomes frequently required to complete tasks reasoning about those interactions during allocation should improve performance. Thus in our experiments we will vary the frequency of required precedence interactions by altering the number of randomly generated debris in the domain. As we increase the number of debris we expect both approaches to suffer in terms of overall reward obtained, but we hope to show that the ATC approach performance degrades at a higher rate than the tiered auction approach.

In the following experiments we tested 7 debris frequency levels: 0, 50, 100, 150, 200, 250, and 300 randomly generated debris. For each debris frequency level we ran both approaches in 25 randomly generated shared instances. In each case we tabulated results after 300 cycles before generating a new domain instance.

Rewards achieved for the two approaches for the 7 debris frequency levels are shown in Figure 7. With 0 debris the two approaches perform identically - in this case, each approach is just bidding using the shortest path distance. Each approach gains less and less of the available reward - which varies around a constant average in each of the trials - as the debris frequency increases. The most important trend to note is that the performance of agents using the tiered auction approach degrades more slowly than the ATC approach, and is significantly higher at all non-zero debris levels. For non-zero debris frequencies, the tiered auction approach outperforms the ATC approach by an average of 76%, with a 140% improvement at the 300 debris level.

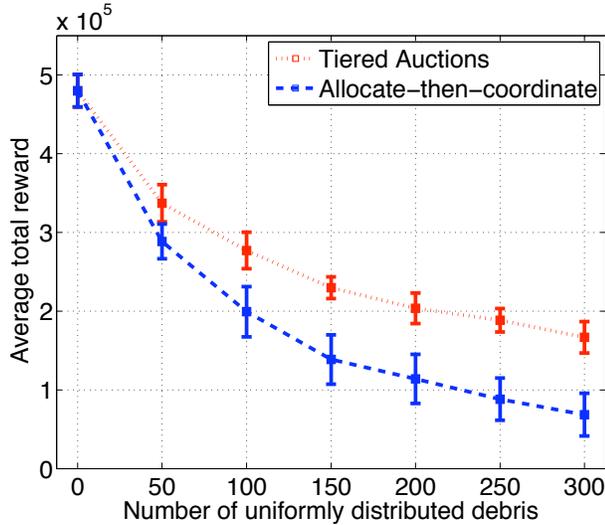


Figure 7: Average performance averaged over 25 trials of 300 time cycles in terms of overall reward for the ATC and tiered auction approaches for each of 7 different levels of debris frequency in a simulated emergency response domain. Standard deviations are shown as error bars.

We can gain more insight into the sources of the performance difference by looking at Figure 8, which shows the average number of completed tasks for each of the two approaches at the different debris levels. We can see that these curves closely resemble those for overall reward. Agents using the tiered auction approach tend to complete tasks more quickly on average as they are not overly optimistic about the time it will take to accomplish tasks during allocation. Considering coordination during allocation allows agents to precisely predict the time it will take to accomplish tasks which lets the system make more informed allocation decisions.

While we have shown that using tiered auctions improves performance over an ATC approach, tiered auctions are relatively computationally expensive. We approximated the computational cost of the approaches by recording wall-clock time for each of algorithms during an auction cycle. An auction cycle consists of all the computation necessary to assign a single task to a single agent - the route planning and sub-auction times for each fire truck and all the bulldozers for each fire being auctioned, and the cost for the winning truck of adopting the fire, which for the ATC approach requires doing full route planning to the allocated fire. This is only a rough approximation of computational cost, as there are many factors that can affect wall clock time. Figure 9 shows the average time per auction for the two different approaches. The time per auction for the ATC approach generally slowly increases with the number of debris but remains very low for all levels, taking less than one-tenth of a second on average. The tiered auction ap-

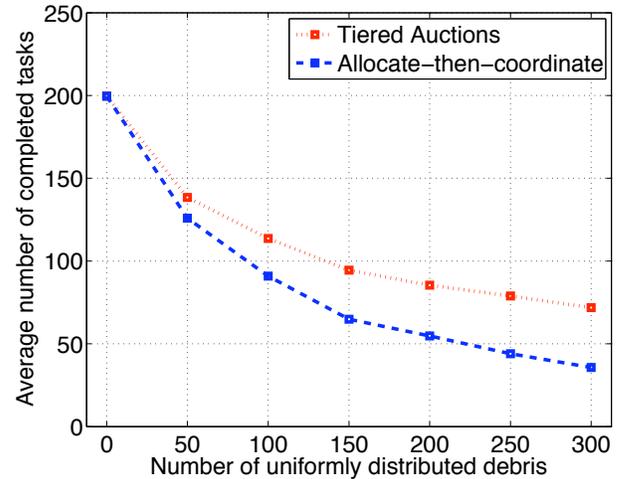


Figure 8: Average number of successfully completed tasks from the trials shown at left.

proach is over 7 times as computationally expensive at the higher debris levels, even without factoring in the communication time that would be required to run sub-auctions on real robots. The substantial increase in performance associated with using tiered auctions comes at the price of added computational costs. We could potentially lower the cost of using tiered auctions by using more restrictive bounding at a cost of reduced performance.

CONCLUSIONS

In this technical report we have detailed methods multi-agent coordination for a precedence-constrained emergency response domain. Our approach uses a novel method, tiered auctions, that equips agents to efficiently search the space of routes and precedence interactions in order to determine plans that will yield high performance. We validated our instantaneous approach by comparison with an “allocate-then-coordinate” approach that used conventional single-tiered auctions and did not allow agents to reason about precedence interactions during allocation; this approach is computationally cheap but results in poor performance, especially for domains with many required precedence interactions.

ACKNOWLEDGEMENTS

This work was sponsored by the U.S. Army Research Laboratory, under contract “Robotics Collaborative Technology Alliance” (contract number DAAD19-01-2-0012). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies or endorsements of the the U.S. Government.

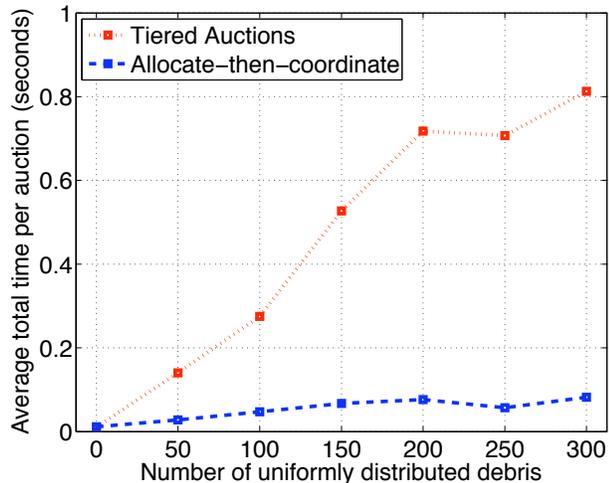


Figure 9: Average time per auction for the ATC and tiered auction approaches for each of 7 different levels of debris frequency in a simulated emergency response domain. Trials were all run in a single thread on a quad-core Xeon 3.8 GHz CPU with 8 GBs of memory.

REFERENCES

- [Berhault et al., 2003] 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 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [Dias, 2004] Dias, M. B. (2004). *TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments*. PhD thesis, Robotics Institute, Carnegie Mellon University.
- [Dias et al., 2004] Dias, M. B., Zlot, R., Zinck, M., Gonzalez, J. P., and Stentz, A. (2004). A versatile implementation of the *TraderBots* approach to multirobot coordination. In *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS)*.
- [Gerkey and Matarić, 2004] Gerkey, B. P. and Matarić, M. J. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9).
- [Gross and Harris, 1998] Gross, D. and Harris, C. M. (1998). *Fundamentals of Queueing Theory*. Wiley, New York, 3rd edition.
- [Jones et al., 2006] Jones, E. G., Browning, B., Dias, M. B., Argall, B., Veloso, M., and Stentz, A. (2006). Dynamically formed heterogeneous robot teams performing tightly-coupled tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- [Jones et al., 2007] Jones, E. G., Dias, M., and Stentz, A. (2007). Learning-enhanced market-based task allocation for oversubscribed domains. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [Koes et al., 2005] Koes, M., Nourbakhsh, I., and Sycara, K. (2005). Heterogeneous multirobot coordination with spatial and temporal constraints. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- [Lagoudakis et al., 2005] Lagoudakis, M., Markakis, E., Kempe, D., Keskinocak, P., Kleywegt, A., Koenig, S., Tovey, C., Meyerson, A., and Jain, S. (2005). Auction-based multi-robot routing. In *Robotics: Science and Systems*.
- [Lagoudakis et al., 2004] Lagoudakis, M. G., Berhault, M., Koenig, S., Keskinocak, P., and Kleywegt, A. J. (2004). Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [Lemaire et al., 2004] Lemaire, T., Alami, R., and Lacroix, S. (2004). A distributed tasks allocation scheme in multi-UAV context. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- [MacKenzie, 2003] MacKenzie, D. C. (2003). Collaborative tasking of tightly constrained multi-robot missions. In *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings of the 2003 International Workshop on Multi-Robot Systems*, volume 2. Kluwer Academic Publishers.
- [Sariel et al., 2007] Sariel, S., Balch, T., and Erdogan, N. (2007). Incremental multi-robot task selection for resource constrained and interrelated tasks. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [Tang and Parker, 2007] Tang, F. and Parker, L. (2007). A complete methodology for generating multi-robot task solutions using AsyMTRe-D and market-based task allocation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- [Vig and Adams, 2006] Vig, L. and Adams, J. A. (2006). Multi-robot coalition formation. In *IEEE Transactions on Robotics*, volume 22 (4).
- [Zlot, 2006] Zlot, R. (2006). *An Auction-Based Approach to Complex Task Allocation for Multirobot Teams*. PhD thesis, Robotics Institute, Carnegie Mellon University.
- [Zlot et al., 2002] Zlot, R., Stentz, A., Dias, M. B., and Thayer, S. (2002). Multi-robot exploration controlled by a market economy. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.