

Agent Coordination with Regret Clearing*

Sven Koenig Xiaoming Zheng
University of Southern California
{skoenig, xiaominz}@usc.edu

Craig Tovey
Georgia Institute of Technology
ctovey@isye.gatech.edu

Richard Borie
University of Alabama
borie@cs.ua.edu

Philip Kilby
National ICT Australia
Philip.Kilby@anu.edu.au

Vangelis Markakis
CWI
V.Markakis@cwi.nl

Pinar Keskinocak
Georgia Institute of Technology
pinar@isye.gatech.edu

Abstract

Sequential single-item auctions can be used for the distributed allocation of tasks to cooperating agents. We study how to improve the team performance of sequential single-item auctions while still controlling the agents in real time. Our idea is to assign that task to agents during the current round whose regret is large, where the regret of a task is defined as the difference of the second-smallest and smallest team costs resulting from assigning the task to the second-best and best agent, respectively. Our experimental results show that sequential single-item auctions with regret clearing indeed result in smaller team costs than standard sequential single-item auctions for three out of four combinations of two different team objectives and two different capacity constraints (including no capacity constraints).

Introduction

We study the distributed allocation of tasks to cooperating agents in real time, where each task has to be assigned to exactly one agent so that the team cost is small or, equivalently, the team performance is high. We do this in the context of multi-robot routing problems, where the agents are robots and the tasks are to visit targets in the plane (Dias et al. 2005). The terrain, the locations of all robots and the locations of all targets are known.¹ Auction-like algorithms (short: auctions) promise to solve multi-robot routing problems with small communication and computation cost since the robots compress information into a small number of bids,

*Some of this research was done while Sven Koenig was on sabbatical at the Australian National University and National ICT Australia. The research was partly supported by NSF awards under contracts ITR/AP0113881, IIS-0098807 and IIS-0350584 as well as seed funding from NASA's Jet Propulsion Laboratory. 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. Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹One can solve multi-robot routing problems in unknown terrain by making assumptions about the unknown terrain, such as the assumption that it is traversable, making it in effect known and thus solvable with auctions. One then runs another auction to re-assign all unvisited targets to robots whenever this assumption turns out to be wrong and thus needs to get revised.

which they compute in parallel and then exchange (Dias et al. 2005). Auctions have been used on actual robots (Gerkey and Mataric 2002; Zlot et al. 2002) and, sometimes in simulation, been applied to sensor networks (Howard and Viguria 2007), mine clearing (Sariel, Balch, and Stack 2006), box pushing (Gerkey and Mataric 2002) and mapping (Simmons et al. 2000). Robotics researchers have recently studied the use of sequential single-item auctions (SSI auctions) for multi-robot routing (Koenig et al. 2006). SSI auctions proceed in several rounds, until all targets have been assigned to robots. During each round, all robots bid on all unassigned targets and the auctioneer then assigns one additional (previously unassigned) target to robots. We study how to improve the team performance of SSI auctions while still controlling the robots in real time by building on algorithmic ideas in the context of vehicle routing (Diana and Dessouky 2004). Our SSI auctions with regret clearing modify the winner determination rule of standard SSI auctions only slightly and leave their bidding rule completely unchanged. Our idea is to assign that target to some robot during the current round whose regret is large, where the regret of a target is defined as the difference of the second-smallest and smallest team costs resulting from assigning the target to the second-best and best robot, respectively.

Multi-Robot Routing

We follow (Koenig et al. 2007) to formalize multi-robot routing problems. A multi-robot routing problem consists of a set of robots $A = \{a_1 \dots a_m\}$ and a set of targets $T = \{t_1 \dots t_n\}$. Any tuple $(T_{a_1} \dots T_{a_m})$ of pairwise disjoint bundles $T_{a_i} \subseteq T$ for all $i = 1 \dots m$ (that is, no target is assigned to more than one robot) is a partial solution of the multi-robot routing problem, with the meaning that robot a_i visits the targets T_{a_i} . Let $c_a(T')$ be the cost needed by robot $a \in A$ to visit the targets $T' \subseteq T$, called **robot cost**, corresponding to the minimal travel distance, fuel consumption or travel time needed to visit the targets from its current location. The travel distances are assumed to satisfy the triangle inequality. The cost of the partial solution, called **team cost**, depends on the team objective. In this paper, we consider two different team objectives, namely the MiniSum team objective and the MiniMax team objective. For the **MiniSum team objective**, the team cost of the partial solution is $\sum_{a \in A} c_a(T_a)$ (that is, the sum of the robot costs), corre-

sponding for example to the fuel consumption of all robots if the robot costs correspond to their individual fuel consumptions for visiting all targets assigned to them, as can for example be important for taking rock probes on the moon. For the **MiniMax team objective**, the team cost of the partial solution is $\max_{a \in A} c_a(T_a)$ (that is, the largest robot cost), corresponding for example to the task-completion time (that is, makespan) if the robot costs correspond to their individual travel times for visiting all targets assigned to them, as can for example be important for search-and-rescue. Any partial solution $(T_{a_1} \dots T_{a_m})$ with $\cup_{a \in A} T_a = T$ (that is, each target is assigned to exactly one robot) is a complete solution of the multi-robot routing problem. We want to find a complete solution of the multi-robot routing problem with a small team cost.

Sequential Single-Item Auctions

Sequential single-item (SSI) auctions solve multi-robot routing problems as follows: Initially, all targets are unassigned. SSI auctions proceed in several rounds, until all targets have been assigned to robots, which then visit the targets assigned to them with minimal travel distance and thus not necessarily in the order in which the targets were assigned to them. During each round, all robots bid on all unassigned targets and the auctioneer then assigns one additional (previously unassigned) target to robots. We now explain the bidding and winner determination rules of standard SSI auctions. Consider any round of a standard SSI auction and assume that robot $a \in A$ has been assigned the targets $T_a \subseteq T$ in previous rounds. Thus, $U = T \setminus \cup_{a \in A} T_a$ is the set of unassigned targets. We leave out the “unassigned” in the following for readability since robots bid only on unassigned targets.

- The **bidding rule** is the following one: Let B be the set of submitted bids. A bid $b \in B$ is a triple (b_a, b_t, b_c) , representing robot b_a , target b_t and bid cost (numerical value of the bid) b_c . For the MiniSum team objective, each robot bids the increase in its robot cost from visiting the target that it bids on in addition to all targets assigned to it in previous rounds, which is similar to previous work on marginal-cost bidding in ContractNet (Sandholm 1996). Formally, $b_c = c_{b_a}(T_{b_a} \cup b_t) - c_{b_a}(T_{b_a})$. For the MiniMax team objective, each robot bids its robot cost of visiting the target that it bids on and all targets assigned to it in previous rounds. Formally, $b_c = c_{b_a}(T_{b_a} \cup b_t)$.
- The **winner determination rule** is the following one: For both the MiniSum and MiniMax team objectives, the auctioneer chooses one of the bids b with minimal bid cost b_c as the winning bid and then assigns the target b_t to robot b_a . Formally, consider any round of a standard SSI auction and let the bid with the smallest bid cost be $b = \arg \min_{b \in B} b_c$. Then, the auctioneer assigns target b_t to robot b_a . Ties can be broken in an arbitrary way.

The bidding and winner determination rules of standard SSI auctions are such that the team cost of the resulting partial solution is as small as possible. Thus, standard SSI auctions make use of a **hill-climbing principle** to achieve

a small team cost. Formally, let $a' \in A$ and $t' \in U$. We define $T_{a',t'}^{a',t'} = T_{a'} \cup t'$ and $T_{a',t'}^{a',t'} = T_{a'}$ for all $a \in A \setminus a'$. For the MiniSum team objective, one can prove that $\sum_{a \in A} c_a(T_a^{b_a, b_t}) = \min_{a' \in A, t' \in U} \sum_{a \in A} c_a(T_a^{a', t'})$ for the winning bid $b \in B$ (Tovey et al. 2005). For the MiniMax team objective, one can prove that $\max_{a \in A} c_a(T_a^{b_a, b_t}) = \min_{a' \in A, t' \in U} \max_{a \in A} c_a(T_a^{a', t'})$ for the winning bid $b \in B$ (Tovey et al. 2005).

Related Work

Standard SSI auctions control robots in real time which is important since robots cannot stop each time they need to assign targets among themselves. We do not expect standard SSI auctions to minimize the team cost since minimizing the team cost for both the MiniSum and MiniMax team objectives is NP-hard (Lagoudakis et al. 2005). However, it is important to achieve a small team cost and thus to decrease the team cost of standard SSI auctions while respecting the real-time constraint. So far, researchers have decreased the team cost of standard SSI auctions by making them more similar to combinatorial auctions (Berhault et al. 2003) while keeping the hill-climbing principle. The idea is that the auctioneer needs to evaluate more complete partial solutions to be able to make good decisions when assigning additional targets to robots:

- **SSI auctions with bundle size $k > 1$** (Zheng, Koenig, and Tovey 2006) proceed in several rounds, until all targets have been assigned to robots. During each round, all robots now bid on sets (called bundles) of at most k targets and the auctioneer then assigns k additional targets to one or more robots. For the MiniSum team objective, each robot bids the increase in its robot cost from visiting the targets in the bundle that it bids on in addition to all targets assigned to it in previous rounds. For the MiniMax team objective, each robot bids its robot cost of visiting the targets in the bundle that it bids on and all targets assigned to it in previous rounds.
- **SSI auctions with rollouts** (Zheng, Koenig, and Tovey 2006) proceed in several rounds, until all targets have been assigned to robots. During each round, all robots bid on all targets and the auctioneer then assigns one additional target to robots. Each robot now bids the team cost of the solution that results if it is assigned the target that it bids on in addition to all targets assigned to it in previous rounds, all other robots are assigned the targets assigned to them in previous rounds, and the resulting partial solution is then (greedily) completed to a solution with a small team cost.

However, the runtimes of these improvements of standard SSI auctions can be large. We therefore explore how to decrease the team cost of standard SSI auctions without increasing their runtime substantially by giving up on the hill-climbing principle.

MiniSum Team Objective

We first study the MiniSum team objective.

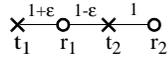


Figure 1: Example 1

Properties of Standard SSI Auctions To determine its bid costs, each robot needs to determine its robot costs, which involves solving an NP-hard traveling salesperson problem (where it does not need to return to its initial location). These calculations can be approximated to run fast. The runtime of standard SSI auctions until all targets are assigned to robots is polynomial if each robot uses the cheapest-insertion heuristic (Lawler et al. 1985) to determine its robot costs approximately. Winner determination is simple and can thus be implemented in a decentralized way without an auctioneer by each robot running the winner determination rule in parallel. The following theorem gives a guarantee on the team cost of standard SSI auctions in form of an upper bound.

Theorem 1 (Lagoudakis et al. 2005) *For the MiniSum team objective, the team cost of standard SSI auctions is at most a factor of two larger than minimal, whether each robot calculates its robot costs exactly or uses the cheapest-insertion heuristic to determine it approximately.*

There is no known instance of multi-robot routing that actually achieves this upper bound. However, the team cost of standard SSI auctions for **Example 1** from Figure 1 is a factor of 1.5 larger than minimal (Tovey et al. 2005). Edges are labeled with their traversal costs, which could be Euclidean planar distances. In the first round of Example 1, robot r_1 bids $1 + \epsilon$ on target t_1 and $1 - \epsilon$ on target t_2 and robot r_2 bids 3 on target t_1 and 1 on target t_2 . Thus, target t_2 gets assigned to robot r_1 . The team cost of the resulting partial solution is indeed as small as possible since assigning target t_1 to robot r_1 results in team cost $1 + \epsilon$, assigning target t_2 to robot r_1 results in team cost $1 - \epsilon$, assigning target t_1 to robot r_2 results in team cost 3 and assigning target t_2 to robot r_2 results in team cost 1. In the second round, robot r_1 bids 2 on target t_1 and robot r_2 bids 3 on target t_1 . Thus, target t_1 gets assigned to robot r_1 . The team cost of the resulting partial solution is indeed as small as possible since assigning target t_1 to robot r_1 results in team cost $3 - \epsilon$ and assigning target t_1 to robot r_2 results in team cost $4 - \epsilon$. To summarize, both targets are assigned to robot r_1 , which results in team cost $3 - \epsilon$ since robot r_1 follows the path r_1, t_2 , and t_1 . This solution does not minimize the team cost since the team cost of assigning target t_1 to robot r_1 and target t_2 to robot r_2 is only $2 + \epsilon$. The ratio $(3 - \epsilon)/(2 + \epsilon)$ approaches 1.5 for small ϵ . Thus, the team cost of standard SSI auctions can be at least a factor of 1.5 larger than minimal.

SSI Auctions with Regret Clearing In the first round of Example 1, target t_2 is assigned to robot r_1 . Yet, this partial solution cannot be completed to a solution with minimal team cost. We now try to understand why standard SSI auctions make this mistake. In the first round, assigning target t_1 to the second-best and best robot, respectively, results in very different team costs (namely, 3 and $1 + \epsilon$) but assigning

target t_2 to the second-best and best robot, respectively, results in similar team costs (namely, 1 and $1 - \epsilon$). Therefore, there is a good chance that target t_2 is assigned to robot r_2 in later rounds if its assignment to robot r_1 is postponed. For example, suppose that target t_1 had been assigned to robot r_1 in the first round. In the second round, assigning target t_2 to robot r_1 would result in team cost $3 - \epsilon$ and assigning target t_2 to robot r_2 would result in team cost $2 + \epsilon$. Thus, target t_2 would be assigned to robot r_2 . In general, later assignments of targets to robots are typically more informed than earlier ones since the partial solutions are more complete then. If a target is assigned to a robot in the current round then one wants to ensure that, if this assignment were postponed, the same assignment would be made in a later round. This is the case if the second-smallest and smallest team costs resulting from assigning the target to the second-best and best robot, respectively, are very different, that is, if their difference is large. We call this difference the **regret** of the target and let the auctioneer assign the target with the largest regret to the robot whose bid on it is lowest. The team cost $\sum_{a \in A} c_a(T_a^{a', t'}) = c_{a'}(T_{a'} \cup t') - c_{a'}(T_{a'}) + \sum_{a \in A} c_a(T_a)$ resulting from assigning target $t' \in U$ to robot $a' \in A$ equals the bid $c_{a'}(T_{a'} \cup t') - c_{a'}(T_{a'})$ of robot a' on target t' plus a constant, namely the team cost $\sum_{a \in A} c_a(T_a)$ before the assignment. Therefore, the difference of the second-smallest and smallest team costs resulting from assigning a target to the second-best and best robot, respectively, and thus the regret of the target equals the difference of the second-smallest and smallest bid on the target. **SSI auctions with regret clearing** thus modify (only) the winner determination rule of standard SSI auctions. They proceed in several rounds, until all targets have been assigned to robots. During each round, the robots bid the increase in their robot cost from visiting the targets that they bid on in addition to all targets assigned to them in previous rounds (as before for the MiniSum team objective) and the auctioneer then assigns one additional target to robots. However, the auctioneer now assigns the target that maximizes the difference of its second-lowest and lowest bids to the robot whose bid on it is lowest. Formally, consider any round of an SSI auction with regret clearing and let the bid with the smallest bid cost on target $t \in U$ be $b^t = \arg \min_{b \in B | b_t = t} b_c$. Then, the auctioneer assigns target $t = \arg \max_{t \in U} ((\min_{b \in B \setminus b^t | b_t = t} b_c) - b_c^t)$ to robot b_a^t . Ties can be broken in an arbitrary way, but we suggest for the auctioneer to consider all targets that maximize the difference of its second-lowest and lowest bids, choose the target from this set with the lowest bid, and assign it to the robot whose bid on it is lowest. SSI auctions with regret clearing then behave like standard SSI auctions in case the regrets of all targets are identical. Winner determination remains simple and can thus again be implemented in a decentralized way without an auctioneer by each robot running the winner determination rule in parallel.

Example of Regret Clearing In the first round of Example 1, robot r_1 bids $1 + \epsilon$ on target t_1 and $1 - \epsilon$ on target t_2 and robot r_2 bids 3 on target t_1 and 1 on target t_2 . The difference of the second-lowest and lowest bids is $2 - \epsilon$ for

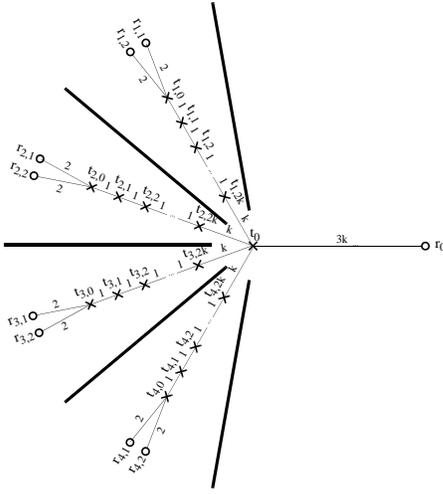


Figure 2: Example 2

target t_1 and ϵ for target t_2 . Thus, target t_1 gets assigned to robot r_1 . In the second round, robot r_1 bids $2 - 2\epsilon$ on target t_2 and robot r_2 bids 1 on target t_2 . Thus, target t_2 gets assigned to robot r_2 . The resulting team cost is minimal.

Properties of Regret Clearing As already discussed, the runtime of standard SSI auctions until all targets are assigned to robots is polynomial if each robot uses the cheapest-insertion heuristic to determine its robot costs approximately. The runtime of SSI auctions with regret clearing is also polynomial under the same condition since only the winner determination rule is different and the new winner determination rule still runs in polynomial time, even though one optimization for standard SSI auctions does not apply to SSI auctions with regret clearing: Each robot needs to submit only its lowest bid during each round for standard SSI auctions since its other bids have no chance of winning. The total number of bids thus equals the number of robots (m) times the number of targets (n). On the other hand, each robot needs to submit a bid on each target for SSI auctions with regret clearing. The total number of bids thus equals $m(n+1)n/2$. This increase in communication cost is unproblematic since each bid can be communicated in a small number of bits.

As already discussed, the team cost of standard SSI auctions is at most a factor of two larger than minimal and a factor of 1.5 larger than minimal for Example 1. We showed that the team cost of SSI auctions with regret clearing is minimal for Example 1. Unfortunately, the team cost of SSI auctions with regret clearing can be larger than the team cost of standard SSI auctions as shown by **Example 2** from Figure 2. The thick lines are walls. Edges are labeled with their traversal costs, which could be Euclidean planar distances. There are $2k+1$ robots called $r_0, r_{i,1}$ and $r_{i,2}$ for $i=1, \dots, k$. There are $2k^2+k+1$ targets called t_0 and $t_{i,j}$ for $i=1, \dots, k$ and $j=0, \dots, 2k$. The auctioneer first assigns target t_0 to robot r_0 and then, one by one, also all other targets to robot r_0 if ties are broken correctly. (The

traversal costs could be changed slightly to achieve this solution no matter how ties are broken, similar to what we did in Figure 1.) The robot pairs $r_{i,1}$ and $r_{i,2}$ for $i=1, \dots, k$ remain unused because they are connected in the same way to the targets. To summarize, all targets are assigned to robot r_0 , which results in team cost $6k^2$ since robot r_0 follows the path $r_0, \{t_0, t_{i,2k}, \dots, t_{i,0}, \dots, t_{i,2k}\}_{i=1}^{k-1}, t_0, t_{k,2k}, \dots, t_{k,0}$. This team cost is not minimal since the team cost of assigning targets $t_{i,0}, \dots, t_{i,2k}$ to robot $r_{i,1}$ for $i=1, \dots, k$ and target t_0 to any robot other than robot r_0 is only $2k^2+3k$ since robot $r_{i,1}$ follows the path $r_{i,1}, t_{i,0}, \dots, t_{i,2k}$ for $i=1, \dots, k$ and one of these robots then visits target t_0 . The ratio $6k^2/(2k^2+3k)$ approaches 3 for large k . Thus, Example 2 shows that the team cost of SSI auctions with regret clearing can be at least a factor of three larger than minimal. The team cost of standard SSI auctions is minimal for Example 2. Thus, SSI auctions with regret clearing do not provide the same good guarantee on the team cost as standard SSI auctions. The following theorem gives a guarantee on the team cost of SSI auctions with regret clearing in form of an upper bound, which might be very weak.

Theorem 2 *For the MiniSum team objective, the team cost of SSI auctions with regret clearing is at most a factor of $2n$ (twice the number of targets) larger than minimal, whether each robot calculates its robot costs exactly or uses the cheapest-insertion heuristic to determine it approximately.*

Proof: We prove the following more general theorem: For the MiniSum team objective, consider an SSI auction where the auctioneer chooses a target according to an arbitrary rule and then assigns it to the robot whose bid on it is lowest. Ties can be broken in an arbitrary way. The team cost of the SSI auction is at most a factor of $2n$ (twice the number of targets) larger than minimal, whether each robot calculates its robot costs exactly or uses the cheapest-insertion heuristic to determine it approximately. For the proof, let $G = (A \cup T, c)$ be the complete graph on the robots and targets, where the weights c correspond to the travel costs. The robot cost $c_a(T)$ then corresponds to the weight of a shortest path that starts at a and visits all vertices in T . The SSI auction proceeds in n rounds. For each round $k=1, \dots, n$, let b^k be the winning bid. Then, $x = \sum_{k=1}^n b^k$ for the team cost x of the SSI auction. We first bound b^k . Let V^k be the set of robots and assigned targets at the beginning of round k , and let \bar{V}^k be the set of unassigned targets at the beginning of round k . V^k and \bar{V}^k define a cut over G with $b^k \in \bar{V}^k$. Consider the cheapest edge that connects b^k and any vertex in V^k and call this vertex v . By the triangle inequality, b^k can be inserted into the robot path that contains v with an increase in robot cost of at most $2c(b^k, v)$. Since the auctioneer assigns b^k to the robot whose bid on it is lowest and the bid costs correspond to these increases in robot cost, it holds that $b^k \leq 2c(b^k, v)$. We now bound $c(b^k, v)$. Let F be a minimum spanning forest of G with roots A and assume that b^k belongs to the tree T_a of F with root a . Since $c(b^k, v)$ is the weight of the cheapest edge that connects b^k and any vertex in V^k and $a \in V^k$, it holds that $c(b^k, v) \leq c(b^k, a)$. By the triangle inequality, $c(b^k, a)$ is no larger than the weight of the path that connects b^k and a in T_a , which is no larger than the weight y of the whole minimum spanning forest. Thus, $c(b^k, a) \leq y$. The robot paths that minimize the team cost form a spanning forest, but not neces-

sarily a minimum spanning forest. Thus, $y \leq z$ for the minimum team cost z . To summarize, $x = \sum_{k=1}^n b_c^k \leq \sum_{k=1}^n 2c(b_t^k, v) \leq \sum_{k=1}^n 2c(b_t^k, a) \leq \sum_{k=1}^n 2y = 2ny \leq 2nz$. ■

MiniMax Team Objective

We now study the MiniMax team objective. For the MiniMax team objective, the team cost of standard SSI auctions is no longer guaranteed to be at most a constant factor larger than minimal, even if each robot calculates its robot costs exactly (Lagoudakis et al. 2005). Furthermore, the experimentally determined average team costs of standard SSI auctions tend to be farther away from minimal for the MiniMax team objective than the MiniSum team objective (Tovey et al. 2005). Therefore, it is even more important to decrease the team cost of standard SSI auctions for the MiniMax team objective than the MiniSum team objective while respecting the real-time constraint. We again let the auctioneer assign the target with the largest regret to the robot whose bid on it is lowest, for the same reason that we gave in the context of the MiniSum team objective. The team cost $\max_{a \in A} c_a(T_a^{a', t'}) = \max(c_{a'}(T_{a'} \cup t'), \max_{a \in A} c_a(T_a))$ resulting from assigning target $t' \in U$ to robot $a' \in A$ equals the maximum of the bid $c_{a'}(T_{a'} \cup t')$ of robot a' on target t' and a constant, namely the team cost $\max_{a \in A} c_a(T_a)$ before the assignment (which the auctioneer knows from the previous round). Therefore, the difference of the second-smallest and smallest team costs resulting from assigning a target to the second-best and best robot, respectively, and thus the regret of the target equals the difference of the second-smallest and smallest bid on the target after all bid costs have been increased to the team cost before the assignment in case they were smaller. We refer to such bids as **increased bids** to distinguish them from the original (unincreased) bids. SSI auctions with regret clearing thus modify again (only) the winner determination rule of standard SSI auctions. They proceed in several rounds, until all targets have been assigned to robots. During each round, the robots bid their robot costs of visiting the targets that they bid on and all targets assigned to them in previous rounds (as before for the MiniMax team objective) and the auctioneer then assigns one additional target to robots. The auctioneer now assigns the target that maximizes the difference of its second-lowest and lowest increased bids to the robot whose original bid on it is lowest. Formally, consider any round of an SSI auction with regret clearing, let the team cost before the assignment be c , and let the increased bid with the smallest bid cost on target $t \in U$ be $b^t = \arg \min_{b \in B \mid b_t=t} \max(b_c, c)$, where b_c denotes the original bid costs. Then, the auctioneer assigns target $t = \arg \max_{t \in U} ((\min_{b \in B \setminus b^t \mid b_t=t} \max(b_c, c)) - \max(b_c^t, c))$ to robot b_a with $b = \arg \min_{b \in B \mid b_t=t} b_c$. Ties can be broken in an arbitrary way, but we suggest that the auctioneer consider all targets that maximize the difference of its second-lowest and lowest increased bids, choose the target from this set with the lowest original bid, and assign it to the robot whose original bid on it is lowest. SSI auctions with regret clearing then behave like standard SSI auctions in case the regrets of all targets are identical. The auctioneer then remembers the team cost $\max(b_c, c)$ and uses it (instead of c) in the next

round as team cost before the assignment.

Experimental Results

SSI auctions with regret clearing solve multi-robot routing problems greedily, just like standard SSI auctions and their variants discussed in the section on related work. We therefore perform experiments to evaluate SSI auctions with regret clearing for multi-robot routing without capacity constraints (where each robot can visit an arbitrary number of targets, as assumed so far) (Dias et al. 2005) and with capacity constraints (where each robot can visit at most a given number of targets, called its capacity, and stops bidding once the number of targets assigned to it equals its capacity) (Koenig et al. 2007). We set all capacities in the latter case to the ratio of the number of targets and robots. We use multi-robot routing on a known eight-neighbor planar grid of size 51×51 with square cells that are either blocked or unblocked. The grid resembles an office environment with walls and doors (Koenig et al. 2007). We average over 25 instances with randomly closed doors for each number of robots and targets. Each robot uses a combination of the two-opt and cheapest-insertion heuristics (Lawler et al. 1985) to determine its robot costs approximately and fast (which usually results in shorter travel distances than the cheapest-insertion heuristic). Space constraints allow us to show only a small number of our results. Table 1 tabulates the team costs (called cost in the table and measured in distance units) and runtimes (called time in the table and measured in seconds). N/A means that the runtime exceeded 500 seconds. The table shows that the runtimes of SSI auctions with rollouts can be so large that they do not respect the real-time constraint, which is why we exclude them from the comparison. SSI auctions with regret clearing tend to run about as fast as standard SSI auctions and faster than SSI auctions with bundle size 3, yet their team costs tend to be smaller for multi-robot routing with capacity constraints (no matter what the team objective is) and multi-robot routing for the MiniMax team objective (no matter whether there are capacity constraints). For the remaining case, namely multi-robot routing without capacity constraints for the MiniSum team objective, standard SSI auctions result in smaller team costs than SSI auctions with regret clearing and also tend to result in smaller team costs than SSI auctions with bundle size 3. To test support for these statements, we perform a statistical analysis of the data. There is no clear way to compare either absolute or relative differences in team costs among scenarios with different parameter settings. We therefore perform a non-parametric analysis, which does not depend on distributional assumptions. A one-sided cumulative binomial test finds the following support for the stated hypotheses:

- For the MiniMax team objective with capacity constraints, SSI auctions with regret clearing are not slower than standard SSI auctions with confidence 0.995, result in smaller team costs than standard SSI auctions with confidence 0.995, are faster than SSI auctions with bundle size 3 with confidence 0.995 and result in smaller team costs than SSI auctions with bundle size 3 with confidence

C	R	T	SSI w. Regret Clearing		Standard SSI		Comparison of Cost		SSI w. Bundle Size 2		SSI w. Bundle Size 3		SSI w. Rollouts		Ideal Hybrid SSI	
			Cost	Time	Cost	Time	Difference	Dominance	Cost	Time	Cost	Time	Cost	Time	Cost	Time
MiniMax Team Objective with Capacity Constraints																
3	8	24	91.06	0.01	111.52	0.01	14.09%	20/25	106.83	0.01	101.03	0.13	63.84	0.24	86.21	0.02
3	12	36	92.87	0.02	109.09	0.02	10.15%	16/25	103.20	0.03	102.51	0.67	58.80	1.75	87.48	0.02
3	16	48	85.44	0.03	109.53	0.03	19.62%	22/25	106.51	0.05	99.88	2.12	62.20	7.53	82.57	0.04
3	20	60	99.20	0.04	117.08	0.04	12.65%	21/25	115.81	0.09	105.97	5.23	60.44	23.66	95.42	0.05
4	6	24	99.52	0.01	121.65	0.01	15.17%	17/25	118.28	0.02	109.39	0.27	66.76	0.39	96.98	0.02
4	9	36	103.58	0.02	125.90	0.02	15.53%	19/25	124.31	0.04	113.82	1.40	66.83	2.88	100.87	0.02
4	12	48	105.33	0.03	126.52	0.03	15.44%	18/25	125.26	0.09	120.68	4.54	68.72	11.99	102.01	0.04
4	15	60	103.63	0.04	121.77	0.04	12.66%	18/25	129.39	0.16	117.30	11.19	61.60	37.72	101.15	0.05
MiniMax Team Objective without Capacity Constraints																
∞	8	24	50.11	0.02	59.45	0.02	14.99%	24/25	60.48	0.03	58.90	0.70	54.84	1.26	50.11	0.02
∞	12	36	42.25	0.03	51.72	0.03	17.45%	22/25	50.54	0.08	50.93	3.78	45.08	10.19	42.13	0.03
∞	16	48	37.00	0.03	45.71	0.03	19.17%	23/25	46.49	0.18	44.99	12.14	39.52	50.44	36.85	0.04
∞	20	60	35.45	0.04	45.12	0.04	21.41%	25/25	44.66	0.36	44.69	31.41	N/A	N/A	35.45	0.05
∞	6	24	62.09	0.02	71.06	0.02	12.11%	21/25	72.21	0.04	68.19	0.88	63.48	1.50	61.93	0.02
∞	9	36	50.80	0.02	64.44	0.03	18.38%	23/25	62.24	0.12	60.87	4.87	52.24	14.76	49.82	0.03
∞	12	48	45.28	0.04	54.33	0.04	15.83%	24/25	53.66	0.30	53.38	16.80	47.76	75.70	45.28	0.05
∞	15	60	41.40	0.05	51.77	0.04	17.96%	23/25	50.92	0.51	50.79	39.54	N/A	N/A	41.16	0.06
MiniSum Team Objective with Capacity Constraints																
3	8	24	363.27	0.01	375.73	0.01	2.61%	15/25	350.18	0.01	357.01	0.13	339.00	0.22	349.63	0.02
3	12	36	464.17	0.02	468.27	0.02	0.29%	11/25	459.41	0.03	453.76	0.66	408.12	1.67	442.74	0.03
3	16	48	553.72	0.03	572.54	0.03	3.06%	17/25	559.04	0.05	557.52	2.13	509.44	7.15	545.19	0.04
3	20	60	635.85	0.04	658.53	0.04	2.26%	14/25	657.19	0.08	644.43	5.22	586.68	22.18	609.09	0.05
4	6	24	346.98	0.01	362.90	0.01	2.95%	14/25	350.05	0.01	351.58	0.24	328.28	0.37	330.63	0.02
4	9	36	431.33	0.02	461.71	0.02	5.65%	19/25	455.18	0.04	446.24	1.21	408.44	2.67	422.77	0.03
4	12	48	542.90	0.03	572.94	0.03	4.24%	11/25	557.39	0.08	549.35	3.97	512.92	11.35	528.96	0.04
4	15	60	579.06	0.04	598.07	0.03	3.23%	17/25	603.03	0.15	581.46	9.67	531.84	34.60	560.63	0.05
MiniSum Team Objective without Capacity Constraints																
∞	8	24	220.29	0.02	215.52	0.02	-2.16%	6/25	217.72	0.06	217.51	1.40	188.44	4.01	214.29	0.02
∞	12	36	244.95	0.04	240.56	0.03	-1.84%	7/25	241.55	0.25	242.57	10.30	210.04	46.33	239.29	0.05
∞	16	48	266.53	0.07	262.68	0.05	-1.52%	5/25	264.34	0.70	261.56	34.10	227.32	208.80	261.42	0.08
∞	20	60	291.58	0.14	284.97	0.07	-2.34%	7/25	286.59	1.27	284.59	88.73	N/A	N/A	283.41	0.18
∞	6	24	239.49	0.02	232.33	0.02	-2.99%	4/25	233.43	0.10	233.25	1.97	203.72	6.40	231.75	0.03
∞	9	36	268.01	0.04	263.75	0.04	-1.61%	7/25	264.53	0.39	264.34	15.37	230.92	68.58	262.14	0.06
∞	12	48	295.38	0.19	288.73	0.07	-2.32%	3/25	290.51	1.40	290.86	78.16	252.16	420.18	288.41	0.23
∞	15	60	309.04	0.15	305.92	0.09	-1.11%	9/25	308.58	2.30	307.32	120.01	N/A	N/A	303.58	0.21

[bold = minimal cost among SSI auctions with regret clearing, standard SSI auctions and SSI auctions with bundle sizes 2 and 3]

Table 1: Experimental Results (C=Capacities, R=Robots, T=Targets)

0.995.

- For the MiniMax team objective without capacity constraints, SSI auctions with regret clearing are not slower than standard SSI auctions with confidence 0.95, result in smaller team costs than standard SSI auctions with confidence 0.995, are faster than SSI auctions with bundle size 3 with confidence 0.995 and result in smaller team costs than SSI auctions with bundle size 3 with confidence 0.995.
- For the MiniSum team objective with capacity constraints, SSI auctions with regret clearing are not slower than standard SSI auctions with confidence 0.95, result in smaller team costs than standard SSI auctions with confidence 0.995, are faster than SSI auctions with bundle size 3 with confidence 0.995 and result in smaller team costs than SSI auctions with bundle size 3 with confidence 0.85 (which is not significant).
- For the MiniSum team objective without capacity constraints, SSI auctions with regret clearing are not faster than standard SSI auctions with confidence 0.995, result in larger team costs than standard SSI auctions with confidence 0.995, are faster than SSI auctions with bundle size 3 with confidence 0.995 and result in larger team costs than SSI auctions with bundle size 3 with confidence 0.995.

We also captured finer-scale data about the team costs of standard SSI auctions and the team costs of SSI auctions with regret clearing. Table 1 tabulates their percent difference (called difference in the table) and the frequency with

which the team cost of SSI auctions with regret clearing is smaller than the team cost of standard SSI auctions (called dominance in the table). A one-sided binomial test with normal approximation ($n = 200$; $p = 0.5$; $\sigma = 7.07$) finds the following support for the stated hypotheses:

- For the MiniMax team objective with capacity constraints, SSI auctions with regret clearing result in smaller team costs than standard SSI auctions with confidence $1 - 10^{-12}$ and a median average difference of 14.6 percent.
- For the MiniMax team objective without capacity constraints, SSI auctions with regret clearing result in smaller team costs than standard SSI auctions with confidence $1 - 10^{-25}$ and a median average difference of 17.7 percent.
- For the MiniSum team objective with capacity constraints, SSI auctions with regret clearing result in smaller team costs than standard SSI auctions with confidence 0.995 and a median average difference of 3.0 percent.
- For the MiniSum team objective without capacity constraints, SSI auctions with regret clearing result in larger team costs than standard SSI auctions with confidence $1 - 10^{-12}$ and a median average difference of 2.0 percent.

The table also contains a column for running both a standard SSI auction and an SSI auction with regret clearing and then using the solution with the smallest team cost, which we call ideal hybrid SSI auctions. (In practice, one would use a classifier to determine whether one expects standard

SSI auctions or SSI auctions with regret clearing to result in smaller team costs and then use the auction recommended by the classifier.) We have argued for Example 1 (and Example 2) under the MiniSum team objective that the team cost of standard SSI auctions is large (and minimal, respectively) but the team cost of SSI auctions with regret clearing is minimal (and large, respectively). Thus, it could be the case that both versions of SSI auctions have complementary strengths. The table shows that ideal hybrid SSI auctions indeed result in smaller team costs than both standard SSI auctions and SSI auctions with regret clearing, although the difference can be small.

Interpretation

To understand the results better, consider a task-allocation problem without synergies among the tasks. The MiniMax team objective with or without capacity constraints then requires one to assign the tasks to robots so that the resulting robot costs are balanced. Standard SSI auctions iteratively assign the task with the smallest cost to a robot. However, it would be more effective to assign the tasks whose costs are large for all robots first and then balance the robot costs using the other tasks. SSI auctions with regret clearing do not necessarily assign the tasks in order of their costs and thus have an advantage over standard SSI auctions. For the MiniSum team objective without capacity constraints, regret is an irrelevant criterion since any task not assigned to a robot in the current round can still be assigned to it in later rounds. Standard SSI auctions use hillclimbing to minimize the team cost directly and thus have an advantage over SSI auctions with regret clearing. For the MiniSum team objective with capacity constraints, regret is a relevant criterion. For example, if the cost of task t_1 is small for robots r_1 and r_2 and the cost of task t_2 is small for robot r_1 but large for robot r_2 , then SSI auctions with regret clearing correctly assign task t_2 to robot r_1 and thus have an advantage over standard SSI auctions.

Conclusions

We studied how to improve the team performance of standard sequential single-item (SSI) auctions while still controlling the robots in real time. Our idea was to assign that target to robots during the current round whose regret is large. Our experimental results show that SSI auctions with regret clearing indeed tend to run about as fast as standard SSI auctions, yet their team costs are smaller for three out of four combinations of two different team objectives and two different capacity constraints (including no capacity constraints). It is future work to combine the various improvements, for example, to study SSI auctions with regret clearing and bundle size $k > 1$.

References

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

Diana, M., and Dessouky, M. 2004. A new regret insertion heuristic for solving large-scale dial-a-ride problems with time windows. *Transportation Research Part B: Methodological* 38(6):539–557.

Dias, M.; Zlot, R.; Kalra, N.; and Stentz, A. 2005. Market-based multirobot coordination: A survey and analysis. Technical Report CMU-RI-TR-05-13, Robotics Institute, Carnegie Mellon University, Pittsburgh (Pennsylvania).

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*.

Koenig, S.; Tovey, C.; Lagoudakis, M.; Markakis, V.; Kempe, D.; Keskinocak, P.; Kleywegt, A.; Meyerson, A.; and Jain, S. 2006. The power of sequential single-item auctions for agent coordination [Nectar Paper]. In *Proceedings of the National Conference on Artificial Intelligence*, 1625–1629.

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.

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.

Sandholm, T. 1996. *Negotiation among Self-Interested Computationally Limited Agents*. Ph.D. Dissertation, Department of Computer Science, University of Massachusetts, Amherst (Massachusetts).

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.

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.

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.; 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.