Multi-Robot Task Allocation with Auctions in Harsh Communication Environments

Michael Otte[1], Michael Kuhlman[2], Donald Sofge[3]

Abstract— **We evaluate three different auction algorithms for multi-robot task allocation when the communication channel is lossy. These include the Sequential Auction, the Parallel Auction, and a generalization of the Prim Allocation Auction called the G-Prim Auction. Each auction is evaluated in two different scenarios: (1) task valuations are random variables drawn from a distribution, and (2) tasks represent locations that must be visited and costs are defined by the extra distance required to visit each location. We derive closed-form solutions for the expected performance of the Sequential Auction and Parallel Auction in Scenario 1, bound the performance of G-Prim in Scenario 1, and bound the performance of the Parallel and Sequential Auctions in Scenario 2.**

I. Introduction

Multi-robot teams often face problems that require dividing a set of tasks among the team's robots, yet multi-robot teams may also operate in environments where communication is unreliable. Communication may be unreliable due to environmental factors such as weather and obstacles, the distance between robots, interference, etc. A popular way to allocate tasks is with an auction. Items are sold to the highest or lowest bidder, where bids are determined by a robot's valuation or cost function, respectively, depending on the problem being solved. Finding an optimal allocation of tasks is equivalent to maximizing the sum of valuations over sold items, or minimizing the sum of costs over sold items. A number of auction algorithms exist, and most are compatible with both the maximum valuation and minimum cost objectives. Auction algorithms differ from each other based on how many rounds are required to sell a particular number of items, how many items each robot can bid for each round, and how many items are sold each round.

We study how communication quality affects the performance of different auction algorithms; in particular, the expected number of tasks won by an agent and the probability an agent does zero tasks. See Figure 1. We assume communication is governed by a Bernoulli process in which each message has an i.i.d. probability p of being sent successfully and a probability $q = 1 - p$ of being dropped, $0 \leq p \leq 1$. Given the Bernoulli communication model, we consider the following two Scenarios:

- S1 Item valuations are random variables drawn from a probability distribution.
- S2 Tasks represent locations, and costs are defined by the extra distance each location adds to a robot's total travel.

For Scenario 1 we derive closed form expressions for the performance of the Parallel and Sequential Auctions as a function of communication quality p, team size, and item count; we prove that the Sequential Auction bounds the performance of a generalization of the Prim Auction [1]. The results for Scenario 1 are then extended to Scenario 2, providing performance equations for the Parallel Auction and performance bounds for the Sequential Auction.

This paper is organized as follows: Related work appears in Section II and preliminaries such as algorithms and notation appear in Section III. The analytical solutions are derived in Section IV, and in Section V they are compared to results from repeated trials in simulation. Discussion and conclusions appear in Sections VI and VII, respectively.

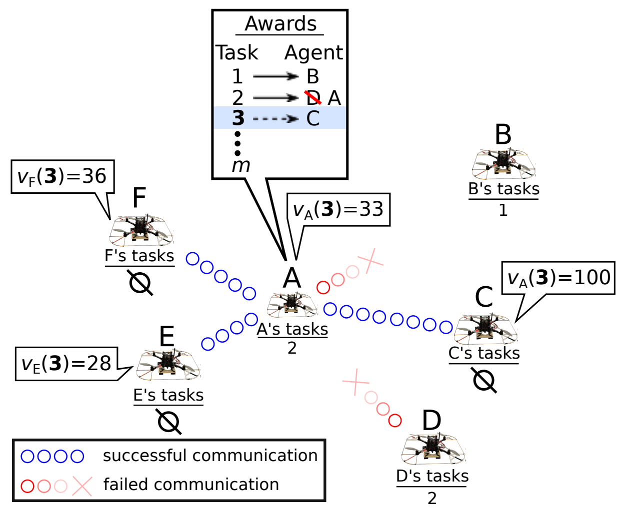

Fig. 1: We study how imperfect communication affects autonomous auctions. In this example, the third of m tasks is advertised and sold to the highest bidder. The message from the auctioneer (A) advertising the sale of task 3 did not reach robots D and B, and they do not bid as a result. Robots A, C, E, and F bid. C wins with the highest bid of $V_C(3) = 100$. A previous message from robot D, acknowledging D's acceptance of task 2, was dropped; thus, the auctioneer must also perform task 2 to ensure that task 2 is completed. Task 1 was successfully awarded to robot B.

This work was performed at the Naval Research Laboratory and was funded by the Office of Naval Research under grant numbers N0001416WX01271 and N0001416WX01272. The views, positions and conclusions expressed herein reflect only the authors' opinions and expressly do not reflect those of the Office of Naval Research, nor those of the Naval Research Laboratory.

[1]National Research Council Postdoc at U.S. Naval Research Lab, Washington D.C. 20375 michael.otte.ctr@nrl.navy.mil

[2]Department of Mechanical Engineering and the Institute for Systems Research, University of Maryland, College Park, MD 20742 mkuhlman@umd.edu

[3]U.S. Naval Research Lab, Washington D.C. 20375 donald.sofge@nrl.navy.mil

II. Related Work

Auction techniques are often used for distributed task allocation in autonomous robotics; surveys can be found in [2] and [3]. We consider the case in which one robot is defined to be the auctioneer *a priori*.

Defining one robot to be the auctioneer is a common approach. That said, other ideas also appear in the literature, including the following: All robots may broadcast their bids so that winners can be calculated in parallel on all robots [4]. Robots may initiate new auctions to re-sell tasks they wish to drop [5]. Robots may share the role of auctioneer [6]–[9]. For other variations the interested reader is referred to [2], [3], since we limit our survey here for brevity.

The auction algorithms that we study include: the Parallel Auction [10], the Sequential Auction [11], and a generalization of Prim Allocation [1] that we call G-Prim. The original Prim Allocation Algorithm assumes tasks are locations, and defines an item's cost as the minimum edge length required to add that item to a spanning tree rooted at the bidding agent. In contrast, G-Prim allows robots to use any valuation or cost function but retains Prim's bidding mechanics: m items are sold over m rounds, and in each round each agent i bids on the single unsold item that i values the most. These mechanics enable G-Prim to have nice properties when communication is poor. All three auctions we study are detailed in Section III.

Previous work on multi-robot task allocation with Sequential Auctions includes: [4], [11]–[14], and with Parallel Auctions includes: [10], [15]. Comparisons of general market based approaches can be found in: [11], [16]–[18]. Previous work evaluating a single auction in interesting communication situations can be found in: [10], [12], [15], [19]–[21]. Our work differs from previous work in that we compare the performance of multiple auction algorithms across a range of communication qualities $p \in [0, 1]$, and we derive expressions for the expected agent participation.

Closely related work includes [21] and [22]. Unreliable communication is assumed in [21], in which Prim Allocation is compared to a distributed technique that uses post-hoc greedy trading for solution improvement; in contrast, we consider different algorithms and valuation functions, and analyze agent participation. Bounds on G-Prim's solution length, when using each of six distance-based heuristics, are reported in [22], assuming perfect communication.

III. Preliminaries

Pseudocode for the Parallel, Sequential, and G-Prim Auctions appears in Algorithms 1, 2, and 3, respectively, assuming a profit *maximization* objective. The alternative cost *minimization* objective can be used by replacing $\arg\max$ by $\arg\min$ and setting unreceived bids to ∞ instead of $-\infty$.

Table I compares auctions based on the number of rounds they require and the number of items bid for and sold each round. Figure 2 shows the phases, messages, and communication graph of a single auction round (identical for all auctions). Each message is either broadcast by the auctioneer or sent once by each non-auctioneer. A broadcast is equivalent to sending the same message once to each agent; we assume the success or failure of a broadcast message reaching different agents is independent. *We assume messages sent from the auctioneer to itself are never dropped.*

TABLE I: High-Level Comparison of Different Auction Types

Auction type	rounds	number of items advertised per round	number of items sold per round
Parallel	1	m	m
Sequential	m	1	1
G-Prim	m	r during round r	1

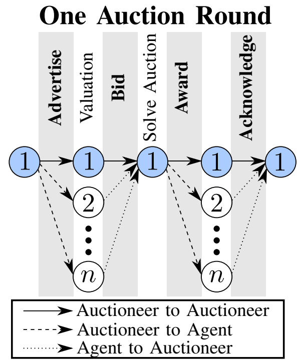

Fig. 2: Phases of an auction round, and the messages sent between them. **Round phases appear boldfaced**, computations during phases appear as plain text. Messages from auctioneer ($a = 1$) to auctioneer are always received, messages from auctioneer to non-auctioneer agents $(2 \dots n)$ and from non-auctioneer agents to auctioneer are sent with probability p and dropped with probability $q = 1 - p$.

Subscripts i, j, and r denote association with a particular item, agent, or round, respectively. There are n agents and m items. a denotes the auctioneer's ID; $a = 1$ by convention. Set Θ contains all items and θ_j is the j-th item. The sets of all sold and unsold items are Θ_{sold} and $\Theta_{\text{unsold}} = \Theta \setminus \Theta_{\text{sold}}$, respectively. During round r, robot i calculates a bid for item θ_j using the valuation function $V_{i,r}(\theta_j)$. Set B_i is agent i's self-maintained set of bids, while $\tilde{B}_i$ is the set of i's bids for which the auctioneer is aware. The set $\tilde{W}_i$ contains the items the auctioneer awards to agent i, while W_i is the set of items robot i knows it has won. Dropped messages may cause $B_i \neq \tilde{B}_i$ and $W_i \neq \tilde{W}_i$. An acknowledgment is denoted C. The cardinality of set '$\cdot$' is denoted $|\cdot|$. The probability of an event '$\cdot$' is $\mathbb{P}(\cdot)$ and the expectation of a value '$\cdot$' is $\mathbb{E}(\cdot)$; probabilities and expectations are defined over the event space of all communication histories and all realizations of valuations (or costs), as defined in Section III-D.

A. Parallel Auction

A Parallel Auction (Algorithm 1) has a single round in which all items are auctioned simultaneously. The auctioneer broadcasts the item list Θ to all robots including itself (line 1). Each robot i calculates a bid for each item θ_j using its valuation function $V_{i,1}(\theta_j)$, and sends its resulting bid list B_i to the auctioneer (lines 3-5). The auctioneer waits for a predetermined length of time to receive bids (lines 6-8), awards each θ_j to the agent that sent the best bid for θ_j (lines 9-11), and then broadcasts the award list $\tilde{W}_{1\dots n}$ (line 12). Winning robots that receive $\tilde{W}_{1\dots n}$ return an acknowledgment C_i (lines 13-15). Finally, the auctioneer takes responsibility for tasks with unacknowledged sales (lines 17-18). Note that while the auctioneer broadcasts $\tilde{W}_{1\dots n}$, each robot i only needs to know the subset $\tilde{W}_i \subset \tilde{W}_{1\dots n}$ (the items that i won).

Algorithm 1 Parallel Auction

On auctioneer a, a.Broadcast(Θ)

On each agent i, **if** i.Receive(Θ) **then**

On each agent i,   **for all** $\theta_j \in \Theta$ **do**

On each agent i, $B_i[j] \leftarrow V_{i,1}(\theta_j)$

On each agent i, i.Send($B_i[1 \ldots m]$)

On auctioneer a, **while** time left **do**

On auctioneer a,   **if** a.Receive($B_i[1 \ldots m]$) **then**

On auctioneer a, $\tilde{B}_{i,1 \ldots m} \leftarrow B_i[1 \ldots m]$

On auctioneer a, **for** $j = 1 \ldots m$ **do**

On auctioneer a, $i \leftarrow \arg\max_i \tilde{B}_{1 \ldots n, m}$

On auctioneer a, $\tilde{W}_i \leftarrow \tilde{W}_i \cup \{\theta_j\}$

On auctioneer a, a.Broadcast($\tilde{W}_{1 \ldots n}$)

On each agent i, **if** i.Receive($\tilde{W}_{1 \ldots n}$) **then**

On each agent i, $W_i \leftarrow \tilde{W}_i$

On each agent i, i.Send(C_i)

On auctioneer a, wait appropriate amount of time

On auctioneer a, **for all** $i \in [1 \ldots n]$ s.t. **not** a.Receive(C_i) **do**

On auctioneer a, $W_a \leftarrow W_a \cup W_i$

Algorithm 2 Sequential Auction

On auctioneer a, $\Theta_{\text{sold}} \leftarrow \emptyset$

On auctioneer a, **while** $\Theta_{\text{sold}} \neq \Theta$ **do**

On auctioneer a,   randomly pick $\theta_j \in \Theta \setminus \Theta_{\text{sold}}$

On auctioneer a, a.Broadcast(θ_j)

On each agent i,   **if** i.Receive(θ_j) **then**

On each agent i, $B_i[j] \leftarrow V_{i,j}(\theta_j)$

On each agent i, i.Send($B_i[j]$)

On auctioneer a,   **while** time left **do**

On auctioneer a,   **if** a.Receive($B_i[j]$) **then**

On auctioneer a, $\tilde{B}_{i,j} \leftarrow B_i[j]$

On auctioneer a, $h \leftarrow \arg\max_i \tilde{B}_{i,j}$

On auctioneer a, $\tilde{W}_h \leftarrow \tilde{W}_h \cup \{\theta_j\}$

On auctioneer a, a.Broadcast(h, j)

On each agent i,   **if** i.Receive(h, j) **and** $i = h$ **then**

On each agent i, $W_i \leftarrow W_i \cup \{\theta_j\}$

On each agent i, i.Send($C_{h,j}$)

On auctioneer a,   wait appropriate amount of time

On auctioneer a,   **if not** a.Receive($C_{h,j}$) **then**

On auctioneer a, $W_a \leftarrow W_a \cup \{\theta_j\}$

On auctioneer a, $\Theta_{\text{sold}} \leftarrow \Theta_{\text{sold}} \cup \{\theta_j\}$

B. Sequential Auction

A Sequential Auction (Algorithm 2) sells 1 item per round for m rounds. During *each round* of a Sequential Auction the auctioneer chooses (e.g., randomly) an unsold item θ_j and sells it using the same phases as a Parallel Auction—advertisement, bidding, winner determination, winner announcement, and acknowledgment.

C. G-Prim Auction

The G-Prim Auction (Algorithm 3) is similar to a Sequential Auction in that one item is sold per round for m rounds; it is also similar to a Parallel Auction in that multiple items are up for sale each round. During round r, each agent bids for the unsold item $\hat{b}_{i,r} \in \Theta_{\text{unsold}}$ that it values the most, and the auctioneer awards the single item that received the best bid (where $\Theta_{\text{unsold}} = \Theta \setminus \Theta_{\text{sold}}$). The number of unsold items $|\Theta_{\text{unsold}}| = m - r + 1$ at the start of round r. Each round has the same phase order as the other two Auctions.

D. Valuation and Cost Scenarios Considered

We now formalize the two Scenarios we consider.

Algorithm 3 G-Prim Auction

for $r = 1 \ldots m$ **do**

On auctioneer a, $\tilde{B}_{i \ldots n, 1 \ldots m} \leftarrow -\infty$

On auctioneer a, $\Theta_r \leftarrow \Theta \setminus \Theta_{\text{sold}}$

On auctioneer a, a.Broadcast(Θ_r)

On each agent i,   **if** i.Receive(Θ_r) **then**

On each agent i,   **for all** $\theta_j \in \Theta_r$ **do**

On each agent i, $B_i[j] \leftarrow V_{i,r}(\theta_j)$

On each agent i, $\hat{b}_{i,j} \leftarrow \arg\max_j B_i[j]$

On each agent i, i.Send($\hat{b}_{i,j}$)

On auctioneer a,   **while** time left **do**

On auctioneer a,   **if** a.Receive($\hat{b}_{i,j}$) **then**

On auctioneer a, $\tilde{B}_{i,j} \leftarrow \hat{b}_{i,j}$

On auctioneer a, $(h, j) \leftarrow \arg\max_{(i,j)} \tilde{B}_{i,j}$

On auctioneer a, $\tilde{W}_h \leftarrow \tilde{W}_h \cup \{\theta_j\}$

On auctioneer a, a.Broadcast(h, j)

On each agent i,   **if** i.Receive(h, j) **and** $i = h$ **then**

On each agent i, $W_i \leftarrow W_i \cup \{\theta_j\}$

On each agent i, i.Send($C_{i,j}$)

On auctioneer a,   wait appropriate amount of time

On auctioneer a,   **if not** a.Receive($C_{i,j}$) **then**

On auctioneer a, $W_a \leftarrow W_a \cup \{\theta_j\}$

On auctioneer a, $\Theta_{\text{sold}} \leftarrow \Theta_{\text{sold}} \cup \{\theta_j\}$

1) Scenario 1, Maximization of Random Valuations: In Scenario 1 we assume all valuations are absolutely continuous random variables drawn i.i.d. at random from the *same* probability density function $f_X(x)$. This is a valid assumption if both: (i) different agents have different preferences and skills that are independently determined, but (ii) the preferences and skills of the agent population as a whole can be described by well behaved distributions. Having all agents draw valuations from $f_X(x)$ guarantees that each possible *ordering* of agent-item valuations is equally likely.

In Scenario 1 valuations are assumed constant over time,

$$V_{i,\hat{r}}(\theta_j) = V_{i,r}(\theta_j) \text{ for all } \hat{r}, r$$

which is a valid assumption if valuations are mutually independent, e.g., there are no synergies between items.

2) Scenario 2, Minimization of Multi-TSP Cost: In Scenario 2 we assume items are locations that must be visited. Cost $V_{i,r}(\theta_j)$ is defined by the extra distance agent i must travel to visit θ_j in addition to the locations $W_{i,r-1}$ that i has already accepted in rounds $1 \ldots r-1$. Let $\ell_{\text{TSP}}(\cdot)$ denote the length of the traveling salesperson (TSP) solution over item set '$\cdot$'. Formally,

$$V_{i,r}(\theta_j) = \ell_{\text{TSP}}(W_{i,r-1} \cup \{\theta_j\}) - \ell_{\text{TSP}}(W_{i,r-1}). \quad (1)$$

In this paper we calculate the *true* multi-TSP length—even for G-Prim. Using the true multi-TSP length is only viable for small numbers of items; the cost metric from Prim Allocation provides a practical alternative for problems involving many items. Prim Allocation (which G-Prim generalizes to any cost or valuation function) was designed especially for Scenario 2 and estimates multi-TSP length using a variant of Christofides TSP approximation algorithm[1] [23].

IV. Analysis

We now analyze the expected performance of the Parallel, Sequential, and G-Prim Auctions in Scenario 1, and then extend this to the expected performance of the Parallel and Sequential Auctions in Scenario 2.

A. Analysis of Parallel Auction in Scenario 1

Communicating a bid to the auctioneer requires receiving the advertisement list and sending a bid message. The probability exactly $k-1$ non-auctioneers communicate a bid to the auctioneer is $p^{2(k-1)}(qp+q)^{n-k}\binom{n-1}{k-1}$. The auctioneer *always* communicates a bid to itself; therefore, when $k-1$ non-auctioneers communicate a bid to the auctioneer the auctioneer receives $(k-1)+1=k$ bids.

The expression $p^{2(k-1)}(qp+q)^{n-k}\binom{n-1}{k-1}$ comes from the facts that: For each of $k-1$ non-auctioneers to submit a bid, advertise messages must be passed successfully from the auctioneer to $k-1$ non-auctioneers (which happens with probability p^{k-1}) and bid messages successfully returned by them (which also happens with probability p^{k-1}). The combined event that the remaining $n-k$ non-auctioneers do not submit bids require that, for each of the $n-k$, *either* (1) an advertise message is successful but the bid message is dropped or (2) the advertise message is dropped; the compound event that either one or the other of these things happen to $n-k$ agents has probability $(qp+q)^{n-k}$. Finally, there are $\binom{n-1}{k-1}$ different ways to assign the non-auctioneers such that $k-1$ non-auctioneers successfully bid and $n-k$ non-auctioneers do not.

Given our assumptions, the probability the auctioneer wins θ_j given k bids are communicated to the auctioneer is $1/k$. Thus, the expected number of items won outright by the auctioneer is:

$$\mathbb{E}\Big(|\tilde{W}_a|\Big) = \sum_{k=1}^{n} \frac{m}{k} p^{2(k-1)}(qp+q)^{n-k}\binom{n-1}{k-1}.$$

The probability a particular $i \neq a$ plus $k-2$ other non-auctioneers communicate a bid message to the auctioneer is $p^{2(k-1)}(qp+q)^{n-k}\binom{n-2}{k-2}$, and so the expected number of items awarded to $i \neq a$ by the auctioneer is:

$$\mathbb{E}\Big(|\tilde{W}_{i\neq a}|\Big) = \sum_{k=2}^{n} \frac{m}{k} p^{2(k-1)}(qp+q)^{n-k}\binom{n-2}{k-2}$$

and the expected number awarded to all non-auctioneers is:

$$\mathbb{E}\Big(\sum_{i\neq a}|\tilde{W}_i|\Big) = (n-1)\mathbb{E}\Big(|\tilde{W}_{i\neq a}|\Big).$$

[1]Prim Allocation (which uses a multi-TSP variant of Christofides TSP approximation algorithm) estimates multi-TSP length by building an incremental spanning forest from agents' locations. A post-processing step is used to obtain a multi-TSP approximation from the spanning forest. This heuristic yields a solution no worse than 1.5 times the optimal length and is calculated in polynomial time. While the true TSP-based solution we consider provides a closer estimate of the optimal multi-TSP path, its runtime is super-polynomial with respect to item number.

Taking responsibility for (or *adopting*) an item requires both winning that item and also receiving the award message. The expected number of items adopted by a single non-auctioneer and the set of all non-auctioneers are, respectively:

$$\mathbb{E}\Big(|W_{i\neq a}|\Big) = p\mathbb{E}\Big(|\tilde{W}_{i\neq a}|\Big)$$

$$\mathbb{E}\Big(\sum_{i\neq a}|W_i|\Big) = (n-1)\mathbb{E}\Big(|W_{i\neq a}|\Big).$$

The expected number of items adopted by the auctioneer includes the items it wins plus all unacknowledged sales:

$$\mathbb{E}(W_{i\neq a}) = \mathbb{E}(\tilde{W}_a) + (pq+q)\mathbb{E}\Big(\sum_{i\neq a}|\tilde{W}_i|\Big).$$

The expected number of items adopted twice, i.e., by the auctioneer as well as a non-auctioneer, is:

$$\mathbb{E}(|W_a \cap (\cup_{i\neq a}W_i)|) = pq\mathbb{E}\Big(\sum_{i\neq a}|\tilde{W}_i|\Big).$$

The probability a non-auctioneer does *not* win item θ_j in a Parallel Auction *assuming its bid is received* is:

$$\mathbb{P}(\theta_j \in \tilde{W}_{i\neq a} \,|\, \zeta) = \textstyle\sum_{k=2}^{n} \frac{k-1}{k} p^{2(k-2)}(qp+q)^{n-k}\binom{n-2}{k-2}$$

where ζ is the event "i's bid is received." Thus, the probability a non-auctioneer wins zero items, assuming ζ, is:

$$\mathbb{P}(\tilde{W}_{i\neq a} = \emptyset \,|\, \zeta) = \mathbb{P}(\theta_j \in \tilde{W}_{i\neq a})^m$$

The probability a non-auctioneer adopts at least one task, i.e., $i \neq a$ wins at least one item and gets the award message, is:

$$\mathbb{P}(W_{i\neq a} \neq \emptyset) = 1 - \Big(q + pq + p^2q + p^3\mathbb{P}(\tilde{W}_{i\neq a} = \emptyset \,|\, \zeta)\Big)$$

B. Analysis of Sequential Auction in Scenario 1

We now switch our focus to the Sequential Auction. The expected number of items won outright by the auctioneer is:

$$\mathbb{E}\Big(|\tilde{W}_a|\Big) = m\sum_{k=1}^{n} \frac{1}{k} p^{2(k-1)}(qp+q)^{n-k}\binom{n-1}{k-1}$$

and the number of items awarded to each non-auctioneer is:

$$\mathbb{E}\Big(|\tilde{W}_{i\neq a}|\Big) = m\sum_{k=2}^{n} \frac{1}{k} p^{2(k-1)}(qp+q)^{n-k}\binom{n-2}{k-2}.$$

These are equivalent to the expectations derived for the Parallel Auction. Thus, all remaining expected values of the Sequential Auction are identical to their Parallel Auction counterparts; we do not repeat them to save space.

The probability that an agent does at least one task in a Parallel Auction is different than in a Sequential Auction, despite the fact that expected number of items adopted by a particular agent is identical for the two auctions. In a Sequential Auction the probability a non-auctioneer $i \neq a$ wins θ_j conditioned on the event ζ_j (its bid is received) is:

$$\mathbb{P}(\theta_j \in \tilde{W}_{i\neq a} \,|\, \zeta_j) = \textstyle\sum_{k=2}^{n} \frac{k-1}{k} p^{2(k-2)}(qp+q)^{n-k}\binom{n-2}{k-2}$$

and the probability a non-auctioneer does at least one task:

$$\mathbb{P}(W_{i\neq a} \neq \emptyset) = 1 - (1 - p\mathbb{P}(\theta_j \in \tilde{W}_{i\neq a} \,|\, \zeta_j))^m.$$

C. Analysis of G-Prim Auction in Scenario 1

We now consider G-Prim. For rounds $r < m$, an agent almost surely gets multiple auction rounds to bid for items it values more than other agents. Let $\neg\zeta_{h,j,r}$ denote the event "agent h's bid for j was not received in round r".

Lemma 1: Given $\theta_j \in \Theta_{\text{unsold}}$ at the beginning of round r and $V_{h,r}(\theta_j) > V_{i,r}(\theta_j)$ for all $i \neq h$ and where $r < m$; then $\mathbb{P}_{r+1}(\theta_j \in \Theta_{\text{unsold}} \,|\, \neg\zeta_{h,j,r}) > 0$.

Proof: By construction $|\Theta_{\text{unsold}}| > 2$ at the beginning of round r when $r < m$. The probability all agents $i \neq h$ bid on some other item $\theta_k \neq \theta_j$ is nonzero and so $\mathbb{P}_r(\theta_j = \hat{b}_{i,j} \,|\, \neg\zeta_{h,j,r}) < 1$ for all $i \neq h$. It follows that $\mathbb{P}_r(\theta_j \in \tilde{B}_{i,j} \,|\, \neg\zeta_{h,j,r}) < 1$ for all $i \neq h$. Finally, if no agents bid on θ_j then θ_j is not sold; i.e., $\theta_j \notin \tilde{B}_{i,j}$ implies that $\theta_j \notin \Theta_{\text{sold}}$ at the end of round r and the beginning of round $r+1$. ■

Corollary 1: For all θ_j such that $h = \arg\max_i V_i(\theta_j)$, $\mathbb{P}_{\text{G-Prim}}(\theta_j \in \tilde{W}_{i\neq a}) > \mathbb{P}_{\text{Sequential}}(\theta_j \in \tilde{W}_{i\neq a})$.

In other words, Corollary 1 states that G-Prim increases $\mathbb{P}(\theta_j \in W_{i\neq a})$ for any θ_j that i values more than any other agent. Theorem 1 leverages Corollary 1 to bound the probability an agent wins zero items in G-Prim based on the Sequential Auction:

Theorem 1: If $m > 1$ and $p < 1$ then:

$$\mathbb{P}_{\text{G-Prim}}(W_{i\neq a} \neq \emptyset) \geq \mathbb{P}_{\text{Sequential}}(W_{i\neq a} \neq \emptyset).$$

Proof: By construction, when $m > 1$ and $n > 1$ there is a greater than $1/n$ probability that an item θ_j exists such that agent h values θ_j more than any other agent; formally, $\mathbb{P}(\exists\theta_j \,|\, h = \arg\max_i V_i(\theta_j)) > 1/n$ for all agents h. Corollary 1 finishes the proof. ■

The effect described in Lemma 1 also increases the expected number of items adopted by non-auctioneers because it reduces the auctioneer's advantage of self communication.

Corollary 2: $\mathbb{E}_{\text{G-Prim}}(|W_{i\neq a}|) \geq \mathbb{E}_{\text{Sequential}}(|W_{i\neq a}|)$.

Because there are only m items, the number of items adopted by the auctioneer must decrease to maintain balance.

Corollary 3: $\mathbb{E}_{\text{G-Prim}}(|W_a|) \leq \mathbb{E}_{\text{Sequential}}(|W_a|)$.

D. Analysis of Scenario 2

Costs in Scenario 2 are defined by the extra TSP length required to visit a new location (Equation 1). When i wins θ_j, the multi-TSP and its sub-length over i's tasks cannot shorten; indeed, it lengthens almost surely[2]. This is formalized in Proposition 1.

Proposition 1: $\mathbb{P}(\ell_{\text{TSP}}(W_{i,r} \cup \{\theta_j\}) > \ell_{\text{TSP}}(W_{i,r})) = 1$.

Lengthening i's multi-TSP path causes i to visit more of the environment (due to the triangle inequality), and decreases i's cost of visiting other locations with a higher probability than it increases it[3]. This is formalized in Proposition 2.

[2]This statement makes the implicit assumption that locations are initially chosen by a random process that would sample the environment densely, in the limit, if the number of locations were allowed to go to infinity. The "almost surely" refers to the fact that, given randomly chosen locations, the chances a new location lies along the old multi-TSP (in which case the multi-TSP length remains the same) is zero.

[3]Note that, this assumes item locations are initially chosen by a random process that would sample the environment densely, in the limit, if the number of locations were allowed to go to infinity.

Proposition 2: Assuming $W_{i,r} = W_{i,r-1} \cup \{\theta_j\}$, then $\ell_{\text{TSP}}(W_{i,r}) > \ell_{\text{TSP}}(W_{i,r-1}) \implies \mathbb{P}(V_{i,r}(\theta_k) < V_{i,r-1}(\theta_k)) > 1/2$, for all $\theta_k \in \Theta_{\text{unsold}}$ at round r.

A lower cost of visiting θ_k increases the chances that an agent will win θ_k. This is formalized in Proposition 3.

Proposition 3: When $p > 0$ and $\theta_k \in \Theta_{\text{unsold}}$ in $r-1$, $V_{h,r}(\theta_k) < V_{i,r-1}(\theta_k) \implies \mathbb{P}(\theta_k \in W_{h,r}) > \mathbb{P}(\theta_k \in W_{i,r-1})$.

Combining Propositions 2 and 3 yields Corollary 4.

Corollary 4: When $p > 0$ and $\theta_j, \theta_k \in \Theta_{\text{unsold}}$ at the beginning of round $r-1$ and $W_{h,r} = W_{h,r-1} \cup \{\theta_j\}$, then $\mathbb{P}(\theta_k \in W_{h,r+1}) > \mathbb{P}(\theta_k \in W_{i,r})$.

Corollary 4 states that if h wins any task θ_j then the probability h wins another task θ_k increases (in Scenario 2). The following Corollary 5 holds because if agent h's chances of winning θ_k increase, the chance that agents $i \neq h$ wins θ_k must decrease.

Corollary 5: When $p > 0$ and $\theta_j, \theta_k \in \Theta_{\text{unsold}}$ in round $r-1$ and $W_{h,r} = W_{h,r-1} \cup \{\theta_j\}$, then $\mathbb{P}(\theta_k \in W_{i\neq h,r+1}) < \mathbb{P}(\theta_k \in W_{i\neq h,r})$.

Given randomly distributed start and item locations, the costs in *round 1* of Scenario 2 meet all assumptions required by the analysis of Scenario 1 (agents draw values from the same distribution and the maximum value objective of Scenario 1 is met by negating Scenario 2's costs). This leads to Proposition 4.

Proposition 4: The probability that i wins round 1 of an auction in Scenario 1 is equal to the probability that i wins Round 1 of the same auction type in Scenario 2.

Corollary 6 holds because Parallel Auctions have one round.

Corollary 6: For the Parallel Auction, all results derived in Scenario 1 are valid for Scenario 2.

We now prove that the equations derived for the Sequential Auction in Scenario 1 become inequalities that provide bounds on the Sequential Auction in Scenario 2 (Lemma 2 and Corollaries 7-10).

Lemma 2: When $p < 1$ and for all θ_j, $\mathbb{P}_{\text{S2, Sequential}}(\theta_j \in \tilde{W}_{a,r}) \geq \mathbb{P}_{\text{S1, Sequential}}(\theta_j \in \tilde{W}_{a,r})$.

Proof: The auctioneer has an advantage over non-auctioneers when $p < 1$ because the auctioneer always has perfect communication with itself. Consequently, a is more likely to win round $r \geq 1$ than $i \neq a$ when $p < 1$; and thus a has an increased chance of winning rounds $r > 1$ by Corollary 4. ■

Lemma 2 has Corollaries 7 and 8, regarding the probability that agents adopt at least one task; and Corollaries 9 and 10, regarding the expected number of tasks adopted by agents.

Corollary 7: When $p < 1$,
$\mathbb{P}_{\text{S2, Sequential}}(\theta_j \in W_{a,r}) \geq \mathbb{P}_{\text{S1, Sequential}}(\theta_j \in W_{a,r})$.

Corollary 8: When $p < 1$,
$\mathbb{P}_{\text{S2, Sequential}}(\theta_j \in W_{i\neq a,r}) \leq \mathbb{P}_{\text{S1, Sequential}}(\theta_j \in W_{i\neq a,r})$.

Corollary 9: When $p < 1$,
$\mathbb{E}_{\text{S2, Sequential}}(|W_a|) \geq \mathbb{E}_{\text{S1, Sequential}}(|W_a|)$.

Corollary 10: When $p < 1$,
$\mathbb{E}_{\text{S2, Sequential}}(|W_{i\neq a}|) \leq \mathbb{E}_{\text{S1, Sequential}}(|W_{i\neq a}|)$,

G-Prim's Scenario 2 performance is similarly bounded by G-Prim's Scenario 1 performance. Formal proofs follow the same logic as for the Sequential Auction, but are less useful

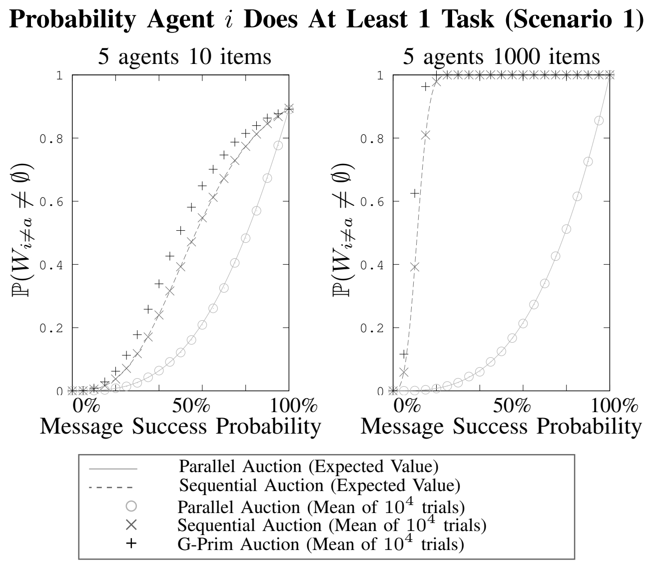

Fig. 3: The probability that a non-auctioneer agent ($i \neq a$) does at least one task, analytical value vs. results from simulations.

because they go in the opposite direction of G-Prim's closed-form bounds for Scenario 1. Nonetheless, averaging over repeated trials of G-Prim in Scenario 1 provides a means of obtaining a numerical bound on its performance in Scenario 2, e.g., as shown in Figure 4-Bottom-Right.

V. Simulations

We run repeated trials in simulation for both Scenario 1 and Scenario 2 across a variety of communication qualities on the range $p \in [0, 1]$. For Scenario 1, every robot determines a unique valuation for each item by drawing a random number from the range $[0, 1]$, and we simulate the two cases where 5 agents divide 10 and 1000 tasks, respectively. For Scenario 2, 5 agents participate in auctions for 10 locations, and start and item locations are drawn uniformly at random from a 100 by 100 kilometer square. The TSP-based costs are recalculated for an agent i after rounds in which i wins an item. We run 10^4 trials per data point and plot the mean values from experiments vs. the expected values predicted by our analysis in Figures 3-5.

VI. Discussion of Results

A. Discussion of Scenario 1

The results from simulations closely match our analysis of Scenario 1 (Figures 3, 4-Top, 4-Middle). The Sequential Auction tends to involve more agents in the final task allocation than the Parallel Auction, despite the fact that the two methods award the same expected number of items to each agent. This happens because, in a Parallel Auction, a single dropped messages from a non-auctioneer will prevent it from bidding on or adopting *any* tasks. In contrast, the Sequential Auction amortizes this risk over m different rounds, each round involving $1/m$ the communication bandwidth of the Parallel Auction. This amortization provides increased benefits as the number of tasks m increases.

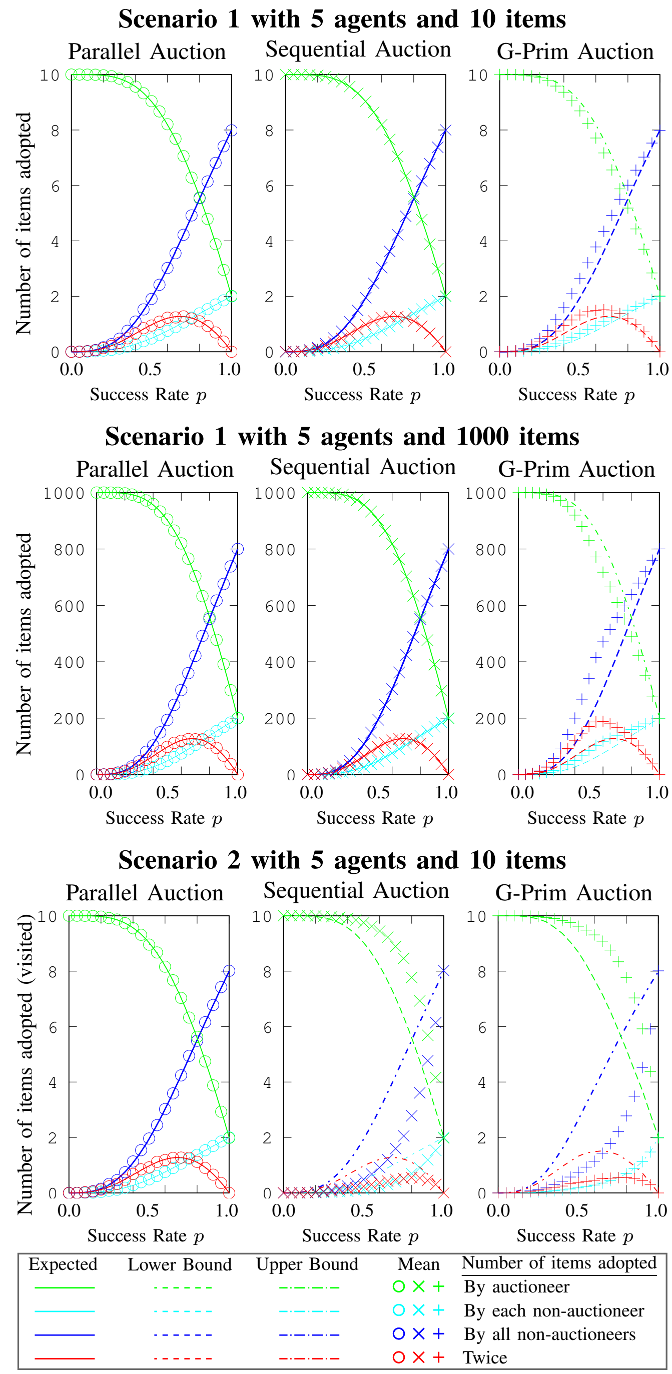

Fig. 4: The number of items adopted by agents over various communication qualities, auctions, and for Scenarios 1 and 2. Note that items are adopted (or visited) twice if an agent receives an award message from the auctioneer but fails to send an acknowledgment message back to the auctioneer (and so the auctioneer also visits the item).

When $p < 1$ the G-Prim Auction enables more agents to win tasks than either the Sequential or Parallel Auctions. By having each agent i bid for the item that i values most, G-Prim reduces the chances that an item highly valued by i is sold to some other agent in the event that a message to/from i is dropped. G-Prim also tends to result in better solutions overall (Figure 5-Right). The price of these advantages is that

Scenario 2 with 5 agents 10 items

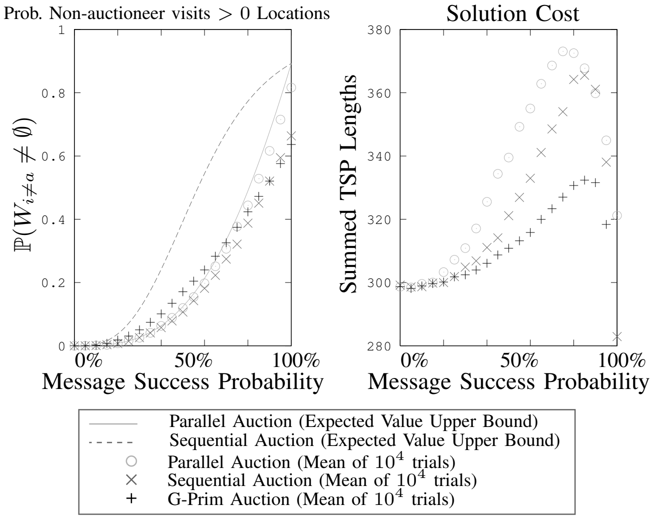

Fig. 5: Left: the probability that a non-auctioneer ($i \neq a$) does at least one task, analytical value vs. results from experiments over various communication qualities for different auction types. Right: Solution quality over various communication qualities.

the auctioneer must use advertisement messages that are $m/2$ times larger than the Sequential Auction, on average.

B. Discussion of Scenario 2

Both simulations and analysis show that the Parallel Auction's performance in Scenario 2 is equivalent to its performance in Scenario 1, and the Sequential Auction's performance in Scenario 2 is bounded by its performance in Scenario 1 (Figures 4-Bottom and 5-Left). In Scenario 2 we observe that, when $p < 1$ for auctions with more than two rounds, non-auctioneers are less likely to win items than in Scenario 1. This happens because winning any item increases the probability of winning additional items in future rounds.

In both Scenarios 1 and 2, communication loss decreases the chances that a non-auctioneer will win an item (and increases the chances that an auctioneer will win an item). As a result better communication correlates with increased agent participation, in general.

VII. Summary and Conclusions

We evaluate the performance of the Parallel Auction, the Sequential Auction, and the G-Prim Auction for multi-robot task allocation in cases where communication between the robots is unreliable (and governed by a Bernoulli process). We derive closed-form solutions for the expected performance of the Sequential and Parallel Auctions and bound the performance of G-Prim in terms of the Sequential Auction's results. We consider two different Scenarios. The first involves maximizing the value of items sold, where item values are random variables. The second assumes items are randomly drawn locations, and defines cost as the extra distance required to visit a location.

The average performance observed over repeated trials in simulation agree with our analysis. When communication is poor in Scenario 1, G-Prim enables more agents to participate in tasks than either the Sequential Auction or the Parallel Auction. When communication is poor in Scenario 2 non-auctioneers are less likely to participate in each task, regardless of auction type.

In general, solution quality is ranked best-to-worst in the order: G-Prim, Sequential, Parallel.

References

[1] M. G. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. J. Kleywegt, "Simple auctions with performance guarantees for multi-robot task allocation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, Sept 2004, pp. 698–705 vol.1.

[2] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, July 2006.

[3] S. Koenig, P. Keskinocak, and C. A. Tovey, "Progress on agent coordination with cooperative auctions." in *AAAI*, vol. 10, 2010, pp. 1713–1717.

[4] D. Vail and M. Veloso, "Dynamic multi-robot coordination," in *In Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II*. Kluwer Academic Publishers, 2003, pp. 87–100.

[5] J. Guerrero and G. Oliver, "Multi-robot task allocation strategies using auction-like mechanisms," *Artificial Research and Development in Frontiers in Artificial Intelligence and Applications*, vol. 100, pp. 111—122, 2003.

[6] R. Smith, "Communication and control in problem solver," *IEEE Transactions on computers*, vol. 29, no. 12, pp. 1104–1113, 1980.

[7] P. Caloud, W. Choi, J.-C. Latombe, C. Le Pape, and M. Yim, "Indoor automation with many mobile robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 1990, pp. 67–72.

[8] S. C. Botelho and R. Alami, "M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement," in *IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 1999, pp. 1234–1239.

[9] M. B. Dias, M. Zinck, R. Zlot, and A. Stentz, "Robust multirobot coordination in dynamic environments," in *IEEE International Conference on Robotics and Automation*, vol. 4. IEEE, 2004, pp. 3435–3442.

[10] M. B. Dias and A. Stentz, "A free market architecture for distributed control of a multirobot system," in *In 6th International Conference on Intelligent Autonomous Systems (IAS-6*, 2000, pp. 115–122.

[11] M. J. Mataric and G. S. Sukhatme, "Task-allocation and coordination of multiple robots for planetary exploration," in *In Proceedings of the 10th International Conference on Advanced Robotics*. Citeseer, 2001.

[12] I. Rekleitis, A. P. New, E. S. Rankin, and H. Choset, "Efficient boustrophedon multi-robot coverage: an algorithmic approach," *Annals of Mathematics and Artificial Intelligence*, vol. 52, no. 2, pp. 109–142, 2008.

[13] C. Pippin and H. Christensen, "A bayesian formulation for auction-based task allocation in heterogeneous multi-agent teams," in *Proceedings of the SPIE*, 2011.

[14] B. P. Gerkey and M. J. Mataric, "Sold!: Auction methods for multi-robot coordination," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002.

[15] R. Zlot, A. Stentz, M. B. Dias, and S. Thayer, "Multi-robot exploration controlled by a market economy," in *IEEE International Conference on Robotics and Automation*, vol. 3, 2002, pp. 3016–3023.

[16] R. C. Cavalcante, T. F. Noronha, and L. Chaimowicz, "Improving combinatorial auctions for multi-robot exploration," in *Advanced Robotics (ICAR), 2013 16th International Conference on*, Nov 2013, pp. 1–6.

[17] E. Schneider, O. Balas, A. T. Ozgelen, E. I. Sklar, and S. Parsons, "Evaluating auction-based task allocation in multi-robot teams," in *AAMAS Workshop on Autonomous Robots and Multirobot Systems (ARMS)*, 2014.

[18] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt, "Robot exploration with combinatorial auctions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2. IEEE, 2003, pp. 1957–1962.

[19] B. J. Moore and K. M. Passino, "Coping with information delays in the assignment of mobile agents to stationary tasks," in *Decision and Control. Proceedings of IEEE conference on*, 2004.

[20] M. Alighanbari and J. P. How, “Decentralized task assignment for unmanned aerial vehicles,” in *Proceedings of the 44th IEEE Conference on Decision and Control*, Dec 2005, pp. 5668–5673.

[21] H. L. Choi, L. Brunet, and J. P. How, “Consensus-based decentralized auctions for robust task allocation,” *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 912–926, Aug 2009.

[22] M. G. Lagoudakis, E. Markakis, D. Kempe, P. Keskinocak, A. J. Kleywegt, S. Koenig, C. A. Tovey, A. Meyerson, and S. Jain, “Auction-based multi-robot routing,” in *Robotics Science and Systems*, 2005.

[23] N. Christofides, “Worst-case analysis of a new heuristic for the travelling salesman problem,” Graduate School of Industrial Administration, Carnegie Mellon University, Technical Report 388, 1976.