

A Dynamic Unit-Demand Auction Supporting Bid Revision^{*}

Chinmayi Krishnappa
Department of Computer Science
University of Texas at Austin
Austin, TX 78712
chinmayi@cs.utexas.edu

C. Greg Plaxton
Department of Computer Science
University of Texas at Austin
Austin, TX 78712
plaxton@cs.utexas.edu

ABSTRACT

We present a dynamic unit-demand auction that supports arbitrary bid revision. Each round of the dynamic auction takes a tentative allocation and pricing as part of the input, and allows each bidder — including a tentatively allocated bidder — to submit an arbitrary unit-demand bid. We establish strong properties of the dynamic auction related to truthfulness and efficiency. Using a certain privacy preservation property of each round of the auction, we show that the overall dynamic auction is highly resistant to shilling. We present a fast algorithm for implementing the proposed auction. Using this algorithm, the amortized cost of processing each bidding operation is upper bounded by the complexity of solving a single-source shortest paths problem on a graph with nonnegative edge weights and a node for each item in the auction. We propose a dynamic price adjustment scheme that discourages sniping by providing incentives to bid early in the auction.

1. INTRODUCTION

Consider the following concrete example of a real-world auction scenario. The developer of a new high-rise condominium project wishes to sell all of its units to the public. In this setting, each bidding agent may assign a different value to each unit, depending on factors such as floor plan, elevation, and view. An agent in this auction is said to have unit-demand preference if the agent is seeking to purchase at most one unit. Agents can have unit-demand preferences in several other settings including auctions for run-of-the-mill items such as bicycles. In a unit-demand auction, the bid of an agent takes the form of a unit-demand preference function: The agent specifies an offer for each of a subset of items, with the understanding that the bid can win at most one item. Typical online auction houses do not support such unit-demand bids. Instead, if many items are to be sold, each is sold in a separate auction. The resulting sequence of single-item auctions forces an

^{*}Supported by NSF Grant CCF-0635203 and THECB NHARP Grant 003658-0235-2007.

agent with unit-demand preferences to guess whether or not to bid on each successive item, since the agent does not know the eventual selling prices of the items. This guesswork degrades the efficiency of the allocation, where the efficiency of an allocation is defined as the sum, over all items v , of the value assigned to v by the agent to which v is allocated. The reason to contemplate selling many items within a single unit-demand auction, or within any combinatorial auction, is to reduce the need for guesswork, thereby enhancing efficiency. By improving efficiency, one has the potential to improve the quality of the outcome for both buyers and sellers alike.

Unit-demand auctions are well understood in the standard sealed-bid framework. In this context, the well-known Vickrey-Clarke-Groves (VCG) [18, 4, 7] mechanism yields a truthful auction that produces an efficient allocation and envy-free pricing [19]. However, the majority of auction sites, including the popular site, eBay, are dynamic. In a dynamic auction, bidding takes place in multiple rounds. In each round, new bid data (bid revision requests and new bids) is received, and an update rule is applied to adjust the tentative outcome (allocation and pricing). The tentative outcome is made public at the end of each round. This dynamic price feedback enables agents to concentrate their value discovery efforts on the most relevant items.

Unit-demand bids are much more expressive than the traditional single item bids and bid formulation is correspondingly more complex. Accordingly, there is a significant chance that a tentatively allocated agent may wish to revise one or more bid components. If a unit-demand auction imposes undue constraints on bid revision, or if the semantics of bid revision introduce additional strategic considerations, then agents may be reluctant to submit unit-demand bids or may only choose to submit bids in the last round of the auction. Such an artificial reduction in the number of bids directly undercuts the main value propositions of dynamic auctions, namely value discovery and improved efficiency.

In this paper, we specify rules for a dynamic unit-demand auction that supports arbitrary bid revision. The reader will note that each round of a dynamic auction is essentially a sealed-bid auction. A guiding principle that we follow in the design of our auction is to use the same sealed-bid auction to resolve each round of the dynamic auction. This guideline is motivated by simplicity of design and ensures that trivial solutions are not considered, e.g, an auction that postpones all of its processing to the last round. In order to motivate the design of

our auction, we analyze the special case of our auction setting for a single item. In what follows, we discuss the design of a dynamic single-item auction supporting arbitrary bid revision. As a natural first approach, we consider using the well-known Vickrey auction to resolve each round. However, it is easy to see that such a dynamic auction discards information on the tentative outcome of each round and is essentially equivalent to running the Vickrey auction exactly once in the last round, after all of the bids have been received. Thus, this approach destroys value discovery, a key feature of dynamic auctions.

Next we consider resolving each round of the dynamic single-item auction in the style of the California auction [17] formulated by Steiglitz. The California auction is a dynamic single-item auction that in each round, allocates the item to the highest bidding agent and posts the second highest bid seen up to that round as the tentative price of the item. The California auction is efficient, satisfies envy-freedom, and retains straightforward bidding in an ex-post Nash equilibrium. The eBay auction is an example of the California auction. We note that each round of the California auction can be viewed as an instance of the Vickrey auction with a reserve price, where in each round, the reserve price of the item is set equal to the tentative price of the item from the previous round. In order to support arbitrary bid revision, in each round, we associate the item with the tentatively allocated agent of the previous round as a reserve agent: if every agent bids less than the tentative price of the item in a round, then the item remains allocated to the reserve agent in that round.

It is straightforward to see that when arbitrary bid revision is allowed, the California auction ceases to be efficient and truthful bidding is no longer an ex-post Nash equilibrium of the auction. For example, consider an instance of the California auction with an item v and agents u_0 and u_1 . Agent u_0 values item v at 20 units in round i and at 10 units in a later round j . Agent u_1 values item v at 19 units in round i and at 20 units in round j . If agent u_1 bids truthfully, then agent u_1 wins item v for 19 units in round j . However, by choosing not to bid in round i and by submitting a bid of 20 units in round j , agent u_1 stands to win item v for a lower price of 10 units in round j . Nonetheless, while the overall auction is not truthful, it can be shown that each round of the auction remains truthful.

Sealed-bid auctions exhibit strong truthfulness and efficiency related properties; yet, the overwhelming majority of online auctions continue to be dynamic. The popularity of dynamic auctions suggests that value discovery is one of the most important requirements of online auctions. As we discussed earlier, value discovery becomes increasingly important with increased bid complexity. From our discussion of the dynamic single-item setting above, it follows that there is an inherent tradeoff between flexibility of bid revision and properties such as truthfulness and efficiency. We find that the strong properties of truthfulness and efficiency are lost when arbitrary bid-revisions are introduced even in the restricted single-item case. We generalize this trade-off to the unit-demand setting in the design of our proposed dynamic auction.

In keeping with our approach of using the same sealed-bid auction to resolve each round of the dynamic auction, we seek

to identify a suitable sealed-bid unit-demand auction that generalizes the California auction to the unit-demand setting. In our discussion of the California auction, we observed that each round of the California auction with arbitrary bid revision is equivalent to an instance of the Vickrey auction where each item is associated with a reserve price and a reserve agent. Such a Vickrey auction with reserves can be viewed as a Vickrey auction in which the item is associated with a put option held by the item's seller, with the reserve agent as the target of the put and the reserve price as the strike price of the put. The put option of the item gives the holder of the put the right to sell the item to the target of the put at the strike price of the put, regardless of market conditions.

In a recent work, we presented a sealed-bid unit-demand auction with put options [11] that generalizes the Vickrey auction with a put option, to the unit-demand case. In the present work, we analyze the dynamic variant of the sealed-bid unit-demand auction with put options. Our proposed dynamic auction proceeds as follows. Each round of our dynamic auction is resolved using an instance of the sealed-bid unit-demand auction with put options. Each item is associated with a put option held by the item's seller. In the first round, the target of each item's put is the seller of the item and the strike price is the reserve price of the item. In each subsequent round, the target and strike price of an item's put are given by the tentatively allocated agent and the tentative price of the item from the previous round.

Since each round of our dynamic auction is resolved using the sealed-bid unit-demand auction with put options, the outcome of each round satisfies all of the equilibrium properties associated with the sealed-bid auction (see Section 4.1). For example, like the sealed-bid auction, each round of our dynamic auction is truthful. In the technical body of the paper, we establish various properties of our dynamic auction that hold over multiple rounds of the auction. In the description of our auction, we resolve each round using the same sealed-bid auction. However, the technical claims in this paper hold more generally for any dynamic auction in which each round is resolved using any sealed-bid auction that satisfies certain subsets of the equilibrium properties detailed in Section 4.1.

With regard to efficiency, recall that in the absence of bid revision, the California auction produces an efficient allocation in each round. In the absence of bid revision, our dynamic auction mimics the behavior of the California auction, and hence achieves the same efficiency guarantee. When agents are allowed to revise their bids in an arbitrary manner, such an efficiency guarantee cannot be achieved without sacrificing other key properties. Since we do not wish to sacrifice these properties, we instead achieve the following relaxed form of efficiency achieved by the sealed-bid unit-demand auction with put options — while the current allocation need not be efficient with respect to the current revision of each bid, it is guaranteed to be efficient with respect to a suitable combination of previous and current revisions. We derive additional efficiency related properties pertaining to multiple rounds of our auction. We show that if an agent is envy-free in a round of the auction, then for any sequence of subsequent rounds in which the agent does not submit a bid revision, the agent continues

to remain envy-free (see Theorem 5.1). We also show that for any sequence of rounds in which an agent is not envy-free and does not submit a bid revision, the auction can only make progress towards achieving efficiency with respect to the most recent bid revision (see Theorem 5.2). We believe that our efficiency-related guarantees are essentially the strongest that can be achieved without sacrificing other properties.

An important consideration in the design on a dynamic auction is its vulnerability to “shill” bidding. If the seller of an item can deduce the maximum price that the agent who is tentatively allocated to the item is willing to pay for the item, then the seller can extract this price without forfeiting sale of the item by submitting a shill offer just below the agent’s offer. Dynamic auctions are known to be particularly susceptible to shill bidding [17]. Thus, a goal of our dynamic auction is to ensure bid privacy for tentatively allocated agents. We establish bid privacy of any tentatively allocated agent u in our dynamic auction with respect to the grand coalition of all agents in the auction except agent u . We assume that in each round of our auction, the grand coalition learns the matching and allocation published in the round, the bid of every agent in the round except the bid of agent u , and whether agent u submitted a bid in the round. We show that our dynamic auction is resistant to shilling by such a grand coalition of agents. Specifically, we show that for any sequence of rounds in which agent u does not submit a bid revision, the grand coalition of agents cannot shill agent u by more than one unit without risking forfeiture of sale in one of the rounds (see Theorem 5.3). Since the running time of our auction is independent of the monetary units used, each unit can be considered to be as low as one cent, thus making our auction highly resistant to shilling.

With respect to scalability, our fast implementation of the proposed dynamic auction (see Section 5.4) processes each bidding operation (i.e., new bid or bid revision) using an amortized constant number of Hungarian [12] augmentations, thereby matching the asymptotic complexity associated with the sealed-bid auction, which uses a single augmentation to process each new bid. The worst-case time complexity of such an augmentation is upper bounded by the cost of running a single-source shortest paths computation on a graph where the number of nodes is proportional to the number of items, and where the number of edges is proportional to the total number of “active” bid components of the tentatively allocated agents. (A component of a unit-demand bid is considered active if the associated offer is at least the current price of the associated item.)

In supporting arbitrary bid revision in the unit-demand setting, our dynamic auction successfully achieves the properties that the California auction achieves with arbitrary revision. For any bid revision operation, our dynamic auction immediately admits a closest approximation to the revised bid, and as prices change over rounds, our auction continually admits closer and closer approximations to the revised bid. An important feature of our auction is that in the special case where bid revisions are consistent – such revisions involve raising all components of the unit-demand bid by the same amount – the outcome of our auction is equivalent the celebrated VCG outcome. For the above mentioned reasons, we believe that our proposed dynamic auction is an appropriate generalization of

the California auction to the unit-demand setting. Due to space constraints, we have chosen to present definitions and lemma statements only, omitting all proofs. Complete proof details are available in our companion technical report [10].

The remainder of the paper is organized as follows. Section 2 discusses related work. Section 3 provides a foundation for the technical presentation to follow. Section 4 reviews the equilibrium properties of the sealed-bid unit-demand auction with put options. Section 5 presents our proposed dynamic auction, establishes various properties of the auction related to truthfulness, efficiency, and shill-resistance, and discusses a fast implementation. Section 7 discusses some extensions and recommendations for the auction.

2. RELATED WORK

Demange et al. [6] present two dynamic unit-demand auctions: an “exact” auction, which we refer to as DGS-exact, and an “approximate” auction, which we refer to as DGS-approximate. In each round, the DGS-exact auction elicits the demand (i.e., set of preferred items at the current prices) of each agent. If there is an overdemand set of items, a minimal overdemand set is found, and the prices of all items in the set are incremented by one. If no overdemand set can be found, the DGS-exact auction terminates and each item is allocated to an agent who demands it. Observe that the DGS-exact auction implicitly supports a limited form of bid revision: An agent is free to revise its unit-demand bid as long as the demands specified in all preceding rounds remain consistent with the revision.

Recognizing the highly restrictive nature of the form of bid revision permitted by the DGS-exact auction, Demange et al. propose the DGS-approximate auction. Like DGS-exact, DGS-approximate is an ascending-price auction. (We remark that Mishra and Parkes [14] describe exact and approximate descending price auctions corresponding to DGS-exact and DGS-approximate.) Agents that are not tentatively allocated are consulted in round-robin order and given the opportunity to either select an item, or pass. If an unallocated agent u selects an item v , the tentative price of item v is increased by a parameter δ , and the tentative allocation is updated to reflect that item v is allocated to agent u . The DGS-approximate algorithm terminates when all of the unallocated agents pass. The DGS-approximate auction has several shortcomings in comparison with our dynamic unit-demand auction: the auctioneer is required to specify a value for the parameter δ ; the outcome is guaranteed to be approximately efficient/truthful, even in the absence of bid revision; there is a tradeoff between the quality of the approximation and the running time of the algorithm; and the bid revision framework is restrictive, since it does not allow for trading of items between tentatively allocated agents.

Gul and Stacchetti [8] present a dynamic auction that generalizes the DGS-exact auction for the setting in which agents demand bundles of items. Gul and Stacchetti show that their auction converges to the smallest Walrasian prices, and that their auction is strategy-proof if the smallest Walrasian prices correspond to the VCG payments. Gul and Stacchetti’s auction, like the DGS-exact auction, supports a limited form of bid revision: An agent is free to revise its bid on a bundle as

long as the demands on the bundle specified in all preceding rounds remain consistent with the revision.

General combinatorial auctions support more complex preferences than unit-demand preferences, such as preferences for bundles of items. Unfortunately, for many combinatorial auctions, the problem of finding an efficient allocation is NP-hard. The computational intractability of general combinatorial auctions motivates the study of specialized combinatorial auctions. Rothkopf et al. discuss special cases (including unit-demand) of combinatorial auctions where the problem of finding an efficient allocation can be solved in polynomial time [15]. Various generalizations of unit-demand have been considered in the literature, including recent work on dynamic auctions for homogeneous [1, 3] and heterogeneous [2, 5, 13] commodities.

3. PRELIMINARIES

In formulating our problem, we make use of some terminology from our recent work on the sealed-bid unit-demand auction with put options [11]. In this section, we review definitions and notation from this work that is relevant to the present paper. We introduce the notions of *bid-graphs* and *configurations*. We use bid-graphs and configurations to model the inputs and outputs of our auction.

3.1 Agents and Items

We refer to the bidders in our auction as agents. In order to break ties among agents, we identify each agent with a binary string identifier. We define the maximum over an empty set of agents as the empty agent ϵ . An item v in our auction is a pair where the first component is a binary string identifier, denoted $id(v)$, and the second component is an integer lower bound on the price of v , denoted $min(v)$. We allow the price of an item in our auction to be negative in order to support procurement-type auctions.

3.2 Bid-Graphs

A *bid-graph* encapsulates a set of items and a set of agents having unit-demand bids on the items. Formally, a bid-graph is an edge-weighted complete bipartite graph $G = (U, V, w)$, where U is a set of agents, V is a set of items, w is a function from the set $U \times V$ to the set of integers, and the following conditions are satisfied: (1) the cardinality of U is at least the cardinality of V ; (2) the empty agent ϵ is not an element of U ; (3) for any pair of distinct items v and v' in V , we have $id(v) \neq id(v')$.

3.3 Configurations

A *configuration* encapsulates a bid-graph along with an associated outcome (allocation and pricing of the items in the bid-graph).

A configuration χ is a triple (G, M, Φ) , where $G = (U, V, w)$ is a bid-graph, M is a maximum cardinality matching (MCM) of G , and Φ is a potential function that maps each item v in V to an integer $\Phi(v)$ such that $\Phi(v) \geq min(v)$. In the definitions that follow, let $\chi = (G, M, \Phi)$ be a configuration where bid-graph $G = (U, V, w)$.

In describing our auction, we start off by assuming that all items are tentatively sold to distinct agents. An item that is not tentatively sold to a bidding agent is assumed to be tentatively sold to a dummy agent representing the seller of the item. Thus, we only consider MCMs in our configurations. By the nature of bidgraphs, the size of an MCM of a given configuration is equal to the number of items in the associated bidgraph.

The function $agents(\chi)$ is the set U and the function $items(\chi)$ is the set V . For any item v in V , we define *potential* (χ, v) as $\Phi(v)$. We define *matched* (χ) as the subset of agents in U that are matched in M , and we define *unmatched* (χ) as the set $U \setminus matched(\chi)$. For any item v in V , we define *match* (χ, v) as the agent u in U such that the edge (u, v) belongs to M . For any agent u in U , we define *gap* (χ, u) as $w(u, v) - \Phi(v)$ if $match(\chi, v) = u$, and as zero otherwise.

We now characterize a suitable directed graph on χ and formulate a reachability condition on this directed graph; we use this reachability condition in describing the solution concept of the sealed-bid unit-demand auction in Section 4.1. We define *digraph* (χ) as the directed graph $(U \cup V, A)$, where A is the set of arcs that includes for each edge (u, v) in $U \times V$ such that $w(u, v) - \Phi(v) \geq 0$ and $w(u, v) - \Phi(v) \geq w(u, v') - \Phi(v')$ for every item v' in $V - v$: (1) an arc (v, u) if edge (u, v) is in M ; (2) an arc (u, v) if edge (u, v) is not in M . For any agent u in $unmatched(\chi)$, we define *items* (χ, u) as the set of items v in V such that there exists a directed path from agent u to item v in *digraph* (χ) . In the terminology of the well-known Hungarian algorithm [12] for weighted bipartite matching, the set *items* (χ, u) is the set of items reachable from agent u in the Hungarian tree rooted at u .

We say that an agent u in U satisfies *envy-freedom* if $gap(\chi, u)$ is nonnegative and $gap(\chi, u) \geq w(u, v) - \Phi(v)$ for all items v in V . We say χ is *Walrasian* if every agent u in U satisfies envy-freedom. We say configuration χ is *semi-Walrasian* if for every agent u in $unmatched(\chi)$ and every item v in *items* (χ, u) , the agent $match(\chi, v)$ satisfies envy-freedom. In the terminology of the Hungarian algorithm, this requirement may be stated more concisely as follows: If an agent u belongs to the Hungarian tree rooted at some non-allocated agent, then u satisfies envy-freedom.

We use configurations to represent both the inputs and outputs of each round of the dynamic auction. If configuration (G, M, Φ) where $G = (U, V, w)$ is the output of a general round i of the auction, then the input of round $i + 1$ is a configuration (G', M, Φ) where G' is a bid-graph of the form (U', V, w') .

4. SEALED-BID UNIT-DEMAND AUCTION WITH PUT OPTIONS

In this section, we review the solution concept of the sealed-bid unit-demand auction with put options [11]. We informally motivate the solution concept as well as provide a formal specification in Section 4.1. The properties of the sealed-bid auction detailed in this section are used to establish strong properties of the proposed dynamic auction in Section 5.

For the classic sealed-bid unit-demand auction, the VCG mechanism returns a Walrasian solution where the pricing is given by the unique minimum price vector over all Walrasian solutions. For the sealed-bid unit-demand auction with put options, only a certain subset of the agents are required to satisfy envy-freedom. Specifically, the outcome is required to be semi-Walrasian (see property 1 in Section 4.1).

A semi-Walrasian configuration χ induces a partition of the items into two sets: the set of items that belong to $items(\chi, u)$ for some agent u in $unmatched(\chi)$, and the remaining items. The items in the former set are said to be *priced at market*, and the remaining items are said to be *priced above market*. For an item v that is priced at market, any positive decrease in the price of v (while leaving the prices of all other items unchanged) yields a solution that is no longer semi-Walrasian. Thus, for an item that is priced at market, the associated put need not be exercised in order to justify the price. For such an item v , the price is required to be at least the strike price (see property 2 in Section 4.1); otherwise, the seller of item v would prefer to exercise the put associated with v . For an item that is priced above market, the price can only be justified via exercise of the associated put; for such an item we require the price to be equal to the strike price (see property 3 in Section 4.1).

The set of items V' priced above market are required to be purchased by the set of agents U' who are targets of the associated puts (see property 4(a) in Section 4.1); the motivation for this requirement is that the items in V' are too expensive to be of interest to any of the remaining agents. The problem of determining a suitable allocation of V' to U' may be viewed as an instance of the house allocation problem [16]; accordingly, standard desiderata is enforced related to envy-freedom (see property 4(b) in Section 4.1) and Pareto-efficiency (see property 5 in Section 4.1).

4.1 Solution concept

Given a configuration $\chi_0 = (G, M_0, \Phi_0)$ as input where $G = (U, V, w)$, the sealed-bid unit-demand auction with put options is a truthful auction whose outcome is a configuration $\chi = (G, M, \Phi)$ satisfying the following conditions which we refer to as properties 1 through 5 for the remainder of this paper:

1. The configuration χ is semi-Walrasian.
2. For any item v in V that is priced at market, $\Phi(v) \geq \Phi_0(v)$.
3. For any item v in V that is priced above market, $\Phi(v) = \Phi_0(v)$.
4. Let V' denote the set of all items in V that are priced above market. Then there is a permutation π of V' such that the following conditions hold.
 - (a) For any item v in V' , $match(\chi, \pi(v))$ is equal to $match(\chi_0, v)$.
 - (b) For any item v in V' having $match(\chi_0, v) = u$, $gap(\chi, u) \geq gap(\chi_0, u)$.
5. For any configuration $\chi' = (G, M', \Phi)$, if there exists an agent u in U such that $gap(\chi, u) < gap(\chi', u)$, then

there exists an agent u' in U such that: (strong version) $gap(\chi', u') < gap(\chi, u')$; (weak version) $gap(\chi', u') \leq gap(\chi, u')$ and u' is matched differently in M and M' .

Note that above conditions are stated in terms of an agent's *gap* rather than the utility. For a unit-demand auction where agents bid truthfully, the *gap* of an agent is equal to its utility, and (the weak version of) Condition 5 corresponds to a solution in the (weak) core. For a truthful auction, a solution in the core satisfies Pareto-efficiency. Our reference to the (weak) core is in the sense defined by Jaramillo and Manjunath [9]; a solution is in the weak core if no subset of agents can exchange their allocated items amongst themselves such that every agent in the subset experiences a strict improvement in utility.

4.1.1 Privacy-preservation

Shill bidding refers to sellers submitting bids on their items with the intent of artificially driving up the item prices. A sealed-bid auction is not vulnerable to shill bidding as each agent submits a single concealed bid and the outcome is computed in a single shot. However, shill bidding is known to be prevalent in dynamic auctions. We identify an additional desired property of the sealed-bid auction that is concerned with preserving the privacy of an agent's bid in the auction. We refer to the following privacy preserving property as property 6 for the remainder of this paper. Property 6 follows from [11, Lemmas 5.28 and 5.29]. In section 5.3.3, we use property 6 to show that our proposed dynamic auction is highly resistant to shilling.

6. Let v be an item priced at market and let $u = match(\chi, v)$. For any integer k such that $gap(shift(\chi, u, k), u) \geq 1$, if $shift(\chi_0, u, k)$ is the input configuration of the round, then the output configuration of the round is given by $shift(\chi, u, k)$.

4.2 Fast implementation

In this section, we briefly review the time bounds of a fast implementation of the sealed-bid unit-demand auction with put options (we refer the reader to [11, Section 5.7] for additional details). In Section 5.4, we use these results to establish a fast amortized time bound for our proposed dynamic auction.

The sealed-bid auction proceeds in two phases. The first phase of the auction corresponds to a proxy-version of the approximate auction proposed by Demange, Gale, and Sotomayor [6]. We say a bid-component is "active" if it is at least equal to the price of the associated item. In the first phase, the bid of each unallocated agent can be processed in time proportional to the time required to solve a single-source shortest paths problem on the active subgraph of the associated bid-graph. Furthermore, the total number of such SSSP computations in any execution of the auction is at most the total number of agents in the auction. The second phase of the auction corresponds to resolving a suitably defined instance of the House Allocation Problem [16] using either the TTC algorithm [16] or the $TC^<$ algorithm [9]. The second phase can be implemented in linear time in the size of the associated bidgraph using the TTC algorithm, and in polynomial time using the $TC^<$ algorithm.

5. DYNAMIC AUCTION

In this section, we present our proposed dynamic unit-demand auction supporting bid revision. The dynamic auction proceeds in rounds and a single application of the sealed-bid unit-demand auction with put options [11] is used to update the tentative allocation and pricing in each round. The output of the last round determines the final allocation and pricing. Below we give an informal description of the input to each application of the sealed-bid auction.

At the beginning of the first round, the tentative pricing is given by the starting prices of the items. Each item v is tentatively allocated to a “dummy agent” for item v whose bid is a single offer on v equal to the reserve price of v . There may be other (non-dummy) agents present in the first round, each of whom has an associated unit-demand bid, which may be arbitrary.

At the beginning of any non-first round, the tentative allocation and pricing is given by the solution to the application of the sealed-bid auction associated with the previous round. The set of agents appearing in the round is equal to the union of the following two sets: (1) agents that were tentatively allocated at the end of the previous round; (2) (non-dummy) agents that were not tentatively allocated at the end of the previous round, and are submitting a new unit-demand bid in the current round. For each agent u in set (1), the associated unit-demand bid in the current round is determined as follows: if u submits a revised bid in this round, then this revised bid is taken to be the bid of u ; otherwise, the bid of u is taken to be the same as in the previous round. We do not allow a dummy agent to revise its bid, since the bid of a dummy agent is merely intended to model the fixed reserve price of the seller.

In Section 5.1, we provide a formal description of the dynamic auction. In Section 5.3, we discuss properties of the dynamic auction related to truthfulness, efficiency, and shill-resistance. In Section 5.4, we discuss an implementation of the dynamic auction with a fast amortized time bound for processing each bidding operation.

5.1 Formal description

The input to the first round of the dynamic auction is a configuration χ satisfying the following conditions: (1) for any item v is $items(\chi)$, the integer $min(v)$ is equal to the seller-specified starting price of item v ; (2) there exists exactly $|items(\chi)|$ agents in $agents(\chi)$ that are designated as dummy agents, and for any dummy agent u and any non-dummy agent u' in the set $agents(\chi)$, we have $u < u'$; (3) for any item v in $items(\chi)$, there is a dummy agent u in $agents(\chi)$ such that $w(u, v)$ is equal to the seller-specified reserve price of v (which is required to be at least the starting price of v), $match(\chi, v) = u$, and $w(u, v') = min(v') - 1$ for any item v' in $items(\chi) - v$.

Each round of the dynamic auction is resolved using the sealed-bid unit-demand auction with put options [11]. We now describe the input of a general non-first round of the auction. Let $\chi = (G, M, \Phi)$ where $G = (U, V, w)$ be the output of round $i - 1$ of the auction. The input to round i is a configuration χ' of the form (G', M, Φ) where $G' = (U', V, w')$ satisfying the following conditions: (1) U does not include an agent u

in $unmatched(\chi)$ if u is either a dummy agent or if u did not submit a bid in round i , (2) For each item v in V and for each agent u that is either a dummy agent or is in $matched(\chi)$ and did not submit a bid in round i , $w'(u, v) = w(u, v)$.

5.2 Auxiliary definitions

A dynamic unit-demand auction D is a sequence of sealed-bid unit-demand auctions where each round of the dynamic auction is resolved using the corresponding sealed-bid auction in the sequence.

For any execution of a dynamic unit-demand auction, we have an associated history of bids that specifies the bids received in each round. We define a bid-history H as a sequence of sets of bids where each set includes a unit-demand bid for each agent in the auction. For any bid-history H , we define $length(H)$ as the length of the sequence H . For any bid-history H and any nonnegative integer $i \leq length(H)$, we define $prefix(H, i)$ as the prefix of H of length i . For any bid-history H , we define $prefix(H)$ as $prefix(H, length(H) - 1)$. For any bid-history H and any agent u , we define $bid(H, u)$ as the bid of agent u in the last set of bids of sequence H . For any bid-history H , any agent u , and any bid β , we define $subst(H, u, \beta)$ as the history H' obtained by substituting $bid(H, u)$ with β .

For any dynamic unit-demand auction D and any bid-history H , we define $config(D, H)$ as the output configuration obtained by running auction D on bid-history H . It follows that for any dynamic unit-demand auction D and any bid-history H , the input and output configurations of each round of auction D can be deduced for the sequence of bids in H .

For any configuration $\chi = (G, M, \Phi)$ where $G = (U, V, w)$, and any agent u in U , we define $envy-free(\chi, u)$ to hold if $gap(\chi, u) \geq 0$ and $gap(\chi, u) \geq w(u, v)$ for any item v in V . For any configuration $\chi = (G, M, \Phi)$ and any agent u such that $\neg envy-free(\chi, u)$, we define $admissible(\chi, u)$ as the set of all bids β in $bids(G)$ such that $envy-free(subst(\chi, u, \beta), u)$. For any bid-history H and any agent u , we say $submit(H, u)$ holds if $bid(H, u) \neq bid(prefix(H), u)$.

5.3 Properties

Recall that a dynamic auction is essentially a sequence of sealed-bid auctions. We say a dynamic unit-demand auction satisfies property 1 if each round of the dynamic auction satisfies property 1 of Section 4.1. We define what it means for a dynamic unit-demand auction to satisfy properties 2, 3, 4, 5, and 6 similarly. In this section, we establish properties of any dynamic unit-demand auction that satisfies certain subsets of properties 1 through 6. Theorems 5.1 and 5.2 establish efficiency-related properties of the dynamic auction and are discussed in Section 5.3.2. Theorem 5.3 establishes a certain shill-resistant property of the dynamic auction and is discussed in Section 5.3.3.

5.3.1 Truthfulness

As we have previously noted in Section 4.1, the sealed-bid unit-demand auction with put options is truthful. Since each round of the dynamic auction is resolved using this sealed-bid auction, it follows that each round of the dynamic auction is truthful.

5.3.2 Efficiency

Each round of the dynamic auction implements the solution concept of Section 4.1. Thus, it follows from property 6 that each round of the dynamic auction produces an outcome that is either (strong version) Pareto-efficient, or (weak version) contained in the weak core.

In Theorems 5.1 and 5.2, we establish efficiency-related properties that hold over multiple rounds of the dynamic auction. We now informally motivate the claims of Theorems 5.1 and 5.2.

Theorem 5.1 establishes that if an agent u is envy-free in a round, then agent u remains envy-free in each subsequent round in which u does not submit a bid.

Consider an agent u who is tentatively allocated to an item v . Assume that agent u submits a bid revision request in round i of the auction, thereby expressing a desire to be allocated to some item v' different from v . After round i , agent u may not be envy-free; informally, this means that the revised bid of u is not fully respected by the auction. Theorem 5.2 establishes that in each round subsequent to round i in which u does not submit a bid revision request and remains allocated to the same item, the dynamic auction makes progress towards respecting the revised bid submitted by u in round i . Specifically, with each successive round, the revised bid of u can only find better and better approximations in the growing set of admissible bids.

LEMMA 5.1. *For any dynamic unit-demand auction D that satisfies properties 1 and 3, any agent u in auction D , and any bid-history H , if u belongs to $\text{unmatched}(\text{config}(D, \text{prefix}(H)))$, then we have $\text{envy-free}(\text{config}(D, H), u)$.*

LEMMA 5.2. *For any dynamic unit-demand auction D that satisfies properties 1, 2, 3, and 4, any bid-history H , and any agent u such that $\text{envy-free}(\text{config}(D, \text{prefix}(H)), u)$, if $\neg \text{submit}(H, u)$, then $\text{envy-free}(\text{config}(D, H), u)$.*

LEMMA 5.3. *For any dynamic unit-demand auction D that satisfies properties 1, 2, and 3, any bid-history H , and any agent u such that $\neg \text{envy-free}(\text{config}(D, H), u)$, if u is matched to the same item v in configurations $\text{config}(D, \text{prefix}(H))$ and $\text{config}(D, H)$, then $\text{admissible}(\text{config}(D, \text{prefix}(H)), u)$ is a subset of $\text{admissible}(\text{config}(D, H), u)$.*

Theorems 5.1 and 5.2 follow directly by induction on Lemmas 5.2 and 5.3 respectively.

THEOREM 5.1. *For any dynamic unit-demand auction D that satisfies properties 1, 2, 3, and 4, any bid-history H , any prefix H' of bid-history H , and any agent u such that $\text{envy-free}(\text{config}(D, H'), u)$, if $\neg \text{submit}(\text{prefix}(H, j), u)$ for $\text{length}(H') < j \leq \text{length}(H)$, then $\text{envy-free}(\text{config}(D, H), u)$.*

THEOREM 5.2. *Let D be a dynamic unit-demand auction that satisfies properties 1, 2, and 3, let H be a bid-history, and*

let H' be a prefix of H . Let u be an agent in D such that $\neg \text{envy-free}(\text{config}(D, H), u)$ and let v be the item such that $\text{match}(\text{config}(D, H), v) = u$. For each j where $\text{length}(H') < j \leq \text{length}(H)$, if $\text{match}(\text{config}(D, \text{prefix}(H, j)), v) = u$ and $\neg \text{submit}(\text{prefix}(H, j), u)$, then $\text{admissible}(\text{config}(D, H'), u) \subseteq \text{admissible}(\text{config}(D, H), u)$.

5.3.3 Shill-resistance

If the seller of an item in a dynamic auction has access to the maximum price that an agent who is tentatively allocated to the item is willing to pay for the item, then the seller can extract this price by submitting a shill offer just below the agent's offer. Thus, a goal of the dynamic auction is to ensure bid privacy for tentatively allocated agents. Below we formalize what it means for an agent to be shilled by Δ units for some non-negative integer Δ . In Theorem 5.3, we establish that no agent in the proposed dynamic auction can be shilled by more than one unit. A consequence of this shill-resistant property is that no seller can artificially raise the price of an item by more than one unit without risking forfeiture of sale. The running time of our auction is independent of the monetary units used; thus each unit can be considered to be as low as one cent, making our auction highly resistant to shilling.

We establish bid privacy of a tentatively allocated agent u in the auction with respect to the grand coalition of all agents in the auction except agent u . For any dynamic unit-demand auction D and any agent u , we define $\text{coalition}(D, u)$ as the set of all agents in D except agent u . The agents in $\text{coalition}(D, u)$ are assumed to learn the following in each round of auction D : (1) the bids of all agents except agent u , (2) whether agent u submitted a bid in the round, (3) the shape (relative differences between offers) of agent u 's bid, and (3) the pricing and allocation at the end of the round. Assumption (1) ensures that our dynamic auction preserves the privacy of agent u even when all of the other agents conspire against u . Assumptions (2) and (3) ensure that our notion of privacy does not merely exploit the fact that agent u is allowed to submit a bid revision in every round. Assumption (4) is natural since the dynamic auction publishes the tentative outcome in every round.

For any dynamic unit-demand auction D , any bid-history H , and any agent u in auction D , we define $\text{possible}(D, H, u)$ as the set of all bids β such that $\text{config}(D, \text{subst}(H, u, \beta)) = \text{subst}(\text{config}(D, H), u, \beta)$. The set $\text{possible}(D, H, u)$ corresponds to the set of possible bids of agent u at the end of the auction that can be deduced by the agents in $\text{coalition}(D, u)$. For any dynamic unit-demand auction D , any bid-history H , any agent u in D , and any bid β in $\text{possible}(D, H, u)$, we define $\text{possible}(D, H, u, \beta)$ as the set of all integers z such that $\text{shift}(\beta, z)$ belongs to $\text{possible}(D, H, u)$.

For any dynamic unit-demand auction D , any bid-history H , and any agent u , we define $\text{risk}(D, H, u)$ to hold if agent u belongs to the set $\text{matched}(\text{config}(D, \text{prefix}(H)))$ and there exists a bid β in $\text{possible}(D, \text{prefix}(H), u)$ such that u belongs to $\text{unmatched}(\text{config}(D, \text{subst}(H, u, \beta)))$.

For any dynamic unit-demand auction D , any bid-history H , and any agent u , we say u is shilled out of Δ units if Δ is the maximum integer such that there exists integers i and $j > i$

that satisfy the following conditions:

- Agent u belongs to $\text{matched}(\text{config}(D, \text{prefix}(H, j)))$
- $\neg \text{submit}(\text{prefix}(H, k), u)$ for each integer k where $i < k \leq j$
- $\neg \text{risk}(D, \text{prefix}(H, k), u)$ for each integer k where $i < k \leq j$
- $\text{gap}(\text{config}(D, \text{prefix}(H, i)), u)$ is greater than or equal to $\text{gap}(\text{config}(D, \text{prefix}(H, j)), u) + \Delta$.

LEMMA 5.4. *For any dynamic unit-demand auction D that satisfies property 1, 2, 3, and 4, any bid-history H , any agent u in $\text{matched}(\text{config}(D, H))$, and any bid β in $\text{possible}(D, H, u)$, either $\text{envy-free}(\text{config}(D, H), u)$, or $\text{shift}(\beta, z)$ belongs to $\text{possible}(D, H, u)$ for any integer z .*

LEMMA 5.5. *For any dynamic unit-demand auction D that satisfies properties 1 and 6, any bid-history H , any agent u in $\text{matched}(\text{config}(D, H))$, and any bid β in $\text{possible}(D, H, u)$, if $\text{envy-free}(\text{config}(D, H), u)$, then there exists a smallest integer z_0 such that for any integer $z \geq z_0$, $\text{shift}(\beta, z)$ belongs to $\text{possible}(D, H, u)$, and either (a) any nonnegative integer is a possible value of $\text{gap}(\text{config}(D, H), u)$, or (b) any positive integer is a possible value of $\text{gap}(\text{config}(D, H), u)$.*

LEMMA 5.6. *For any dynamic unit-demand auction D that satisfies properties 1, 2, 3, 4, and 6, any bid-history H , and any agent u in $\text{matched}(\text{config}(D, H))$, if $\neg \text{submit}(H, u)$ and $\neg \text{risk}(D, H, u)$, then for any bid β in $\text{possible}(D, H, u)$, we have $\text{possible}(D, H, u, \beta) = \text{possible}(D, \text{prefix}(H), u, \beta)$.*

THEOREM 5.3. *For any dynamic unit-demand auction D that satisfies properties 1, 2, 3, 4, and 6, and any bid-history H , no agent in auction D can be shilled by more than one unit.*

5.4 Scalability

In this section, we briefly sketch the details of a fast implementation of the dynamic auction. In each round of the dynamic auction, new bid data (i.e., bid revision requests from tentatively allocated agents, and bids from unallocated agents) is received and the tentative allocation and pricing is updated using an instance of the sealed-bid unit-demand auction with put options [11]. Recall from Section 4.2 that in the first round of the sealed-bid auction, the bid of each unallocated agent can be processed in time proportional to the time required to solve a single-source shortest paths (SSSP) problem on the active subgraph of the associated bid-graph. At the start of a round, one or more tentatively allocated agents are not envy-free. If a tentatively allocated agent u who is not envy-free becomes unallocated in the round, then u is added to the set of unallocated agents whose bids are yet to be processed. For each such agent u , the number of tentatively allocated agents who are not envy-free decreases by at least one as the only way a tentatively allocated agent can cease to be envy-free is by revising its bid. When we use an SSSP computation to process

the bid of the now-unallocated agent u , we can charge the cost of this SSSP computation to the most recent bid revision of u . Consequently, the total number of SSSP computations performed across all rounds is at most the total number of bidding operations (i.e., bid revisions or new bids) over all rounds. For our fast implementation, we choose the TTC [16] algorithm to implement the second phase of the sealed-bid auction used in each round. From Section 4.2, the TTC algorithm can be implemented in time linear in the size of the active subgraph.

In summary, it is possible to implement the dynamic auction in such a way that the amortized cost of each bidding operation is close to linear in the size of the active subgraph of the bid-graph, which is at most quadratic in the number of items. Moreover, in many practical auction settings, the average number of active bid components of a tentatively allocated agent is likely to be small, say at most a constant. In such settings, the number of active bid components is linear in the number of items, and hence the amortized cost of each bidding operation is close to linear in the number of items.

6. SNIPING FEES

In this section we describe how to modify our dynamic unit-demand auction to encourage early bidding, while preserving all of the theoretical properties established earlier in the paper. In Section 6.1 we review a standard technique for incorporating agent-specific adjustments into the selling prices of the items. In Section 6.2, we generalize this technique to allow for price adjustments that may increase from one round to the next. Such dynamic price adjustments are used to discourage “sniping”, i.e., waiting until close to the end of the auction to submit a bid. Sniping diminishes agents’ ability to focus value discovery efforts on the most relevant items, thereby increasing participation costs and degrading efficiency.

6.1 Static Price Adjustments

At the conclusion of a typical single-item online auction, the item is shipped to the winning agent. The winning agent pays the auction price plus shipping costs. The cost of shipping is agent-specific, in general, because it may vary depending on the shipping address. The cost of shipping is typically made available to the agents during the auction (e.g., via a shipping calculator). Viewed more abstractly, the seller publishes a static function $\text{adjustment}(u)$ as part of the auction listing, and if agent u wins the auction, then agent u pays the auction price plus $\text{adjustment}(u)$.

Such an abstraction is also useful for selling an item with multiple “variants”. For example, consider a computer that can be sold with or without a monitor. The auction listing for the computer might specify an additional charge for the optional monitor. Such variant-related charges might be agent-specific (e.g., due to the cost of shipping the monitor), and in general might be positive or negative. The auction listing of the seller provides the necessary information (e.g., shipping calculator, fixed price adjustments for different variants) to allow each agent u to determine the relevant price adjustment $\text{adjustment}(u)$ to be paid in the event that agent u wins the auction. We view the adjustment as a function of the agent only, as opposed to the agent and variant, because the agent selects the relevant variant based on the published cost adjustments.

In this sense, the problem of supporting multiple variants of an item is reduced to the single-variant case.

The framework discussed above generalizes immediately to the unit-demand setting, where we have a static price adjustment function $adjustment(u, v)$ that specifies the amount to be added to the auction price to determine the total price paid by agent u for item v . It is natural to ask whether the theoretical properties established for our auction continue to hold in the presence of such an adjustment function. Apart from the price adjustment performed at the end of the auction, the computations performed by our auction depend only on the non-adjusted bids. Consequently, it is straightforward to argue that all of the theoretical properties established for our auction continue to hold with respect to the non-adjusted bids/prices. (Regarding our claim that each individual round of our auction is truthful, we point out that a non-adjusted bid of agent u is truthful if the corresponding adjusted bid is equal to the truthful preferences of agent u .)

6.2 Dynamic Price Adjustments

The static price adjustment framework discussed in Section 6.1 reflects standard practice in single-item auctions, and generalizes straightforwardly to the unit-demand setting. We now introduce a variation of this framework in which there is a separate price adjustment function $adjustment_i$ for each round i of the auction. We require that the choice of the function $adjustment_i$ is determined by the public component of the bidding history up to the start of round i , and that for any agent u , item v , and rounds i and j such that $i < j$, we have $adjustment_i(u, v) \leq adjustment_j(u, v)$.

When an agent u wins an item v , agent u pays the auction price plus $adjustment_i(u, v)$, where i is the index of the earliest round such that for all rounds with index at least i , the output configuration $\chi = (G, M, \Phi)$ satisfies $w(u, v) - \Phi(v) \geq \min(0, gap(\chi, u))$. Roughly speaking, the latter condition checks whether agent u 's unit-demand bid still has the possibility of winning item v in a later round, even if it is left unchanged.

Reasoning as in the case of a static price adjustment function discussed in Section 6.1, it is straightforward to argue that our dynamic price adjustment scheme continues to enjoy all of the theoretical properties established in earlier sections of the paper. In this regard, we remark that our requirement that function $adjustment_i$ is determined by the public component of the bidding history up to the start of round i ensures that the shill-resistance of the auction is preserved. (In the absence of this requirement, the choice of the function $adjustment_i$ could reveal private information related to the bids of the tentatively allocated agents.) The scalability of our auction is unaffected since it is easy to compute the price adjustments to be applied at the end of the auction. Indeed, we can compute tentative price adjustments at the end of each round without increasing the asymptotic complexity of processing a round.

The requirement that $adjustment_i(u, v)$ is nondecreasing in i is motivated by our desire to encourage early bidding in the auction. Suppose agent u wins item v , and our price adjustment rule prescribes that agent u pays the auction price plus $adjustment_i(u, v)$. We view the nonnegative value obtained by

subtracting $adjustment_1(u, v)$ from $adjustment_i(u, v)$ as the “sniping fee” incurred by agent u for not bidding earlier. (Remark: The term “sniping” is often used narrowly to refer to submitting a bid in the last few seconds of an auction. Here we use the term more broadly, since our sniping fee structure can be multi-tiered to discriminate between bids submitted with arbitrarily varying amounts of time remaining in the auction.)

We now describe a simple but practically important use case of the dynamic price adjustment framework introduced above. Consider a unit-demand auction in which the listing of each item v specifies — at the outset of the auction — the value of $adjustment_i(u, v)$ for all agents u and rounds i . For $i = 1$, these values can be used to model shipping costs and item variants as discussed in Section 6.1. The sniping fee applicable to bids submitted in the first round is zero. For any agent u and round $i > 1$, the quantity $adjustment_i(u, v) - adjustment_{i-1}(u, v)$ models the nonnegative change in sniping fee from round $i - 1$ to round i . A potential drawback of this scheme is that the sniping fees for an item v accumulate even while the item remains tentatively allocated to the dummy agent for item v (because the reserve price has not been reached). It may be preferable for the seller of item v to specify how sniping fees are to grow once the reserve price has been met. Such a more complex sniping fee schedule — which has a nontrivial dependence on the bidding history — still falls well within the general dynamic price adjustment framework introduced above.

We now discuss considerations associated with the design of a suitable sniping fee schedule for a given item in the auction. For the sake of concreteness, we pursue this discussion in the context of a specific popular online auction format: A continuous auction with a fixed one-week duration. (In a continuous auction, the tentative pricing and allocation is updated immediately once a bidding operation is received. Equivalently, we can imagine that each round corresponds to a fixed, infinitesimally small, interval of time.) Similar considerations arise in the design of sniping fee schedules for other auction formats.

Under one simple sniping fee schedule for a one-week continuous auction, the sniping fee might increase linearly from zero — at the time when the reserve is met — to a seller-specified maximum value at the end of the auction. However, such a schedule is unlikely to be appropriate. Notice, for example, that the sniping fee would remain virtually constant over the last hour of the auction. From the perspective of allowing competing agents to engage in additional value discovery in response to one’s bid, there is a significant difference between bidding with ten seconds left in the auction, with one minute left, with five minutes left, or with an hour left. This observation leads us to conclude that the sniping fee should increase more rapidly as the time remaining in the auction diminishes. For example, it might be reasonable to make the sniping fee proportional to the logarithm of the ratio of the auction duration to the time remaining. Doing so results in an additive increase in the sniping fee whenever the time remaining decreases geometrically, and a geometric decrease in the time remaining has a qualitative impact on the ability of agents to engage in value discovery.

7. DISCUSSION

In our current presentation of the auction, we have assumed that the set of items for sale in the auction is static. It is straightforward to modify the auction to allow new items to be introduced in each round. Further, we have assumed that all items in the auction have the same expiry time. It is possible to relax this assumption. For example, we can specify a separate expiration time for each item in the auction, and allow unit-demand bidding across items that expire within the same interval of time (e.g., the same day).

The binary string identifiers associated with agents and items in our dynamic auction provide fixed total orders over the sets of agents and items. The properties of our auction continue to hold even if these total orders are changed from one round of the auction to the next. We recommend using a single total order over the items for all rounds of the dynamic auction; for example, this total order could be derived by sorting a fixed set of item identifiers. We recommend a slightly more complex scheme for determining the total order over the agents to be used in each round. First, we recommend that all dummy agents be ordered lower than all non-dummy agents in every round; this ensures that an item can be sold to a non-dummy agent at the starting price. Within the set of dummy agents, we recommend using the same (arbitrary) fixed total order in all rounds.

Within the set of non-dummy agents, we recommend using a dynamic timestamp-based ordering, where the timestamp of an agent is determined as follows. In the first round, all agents are assigned the same timestamp. In any non-first round i , recall that the agents may be partitioned into sets (1) agents that were tentatively allocated at the end of the previous round; and (2) (non-dummy) agents that were not tentatively allocated at the end of the previous round, and are submitting a new unit-demand bid in the current round. Timestamp i is assigned to all of the agents in set (2). Each agent u in set (1) is assigned the minimum timestamp j less than i such that u is allocated in the solution associated with all applications of the sealed-bid unit-demand auction with put options in rounds j through $i - 1$. Having determined these timestamps, we recommend ordering any pair of agents u and u' participating in round i as follows: if u and u' have distinct timestamps, then the agent with the higher timestamp is considered lower; if u and u' have equal timestamps, then the order of the agents is determined by an arbitrary fixed total order. The motivation for the proposed dynamic timestamp-based scheme is that it breaks ties in favor of agents who have been allocated longer.

The single-item auction mechanism employed by eBay is essentially a dynamic second-price auction. We have shown how to generalize this popular auction format to the unit-demand case, while supporting arbitrary bid revision by tentatively allocated bidders. Our auction maintains strong properties related to efficiency, truthfulness, privacy preservation, and scalability. We have implemented our auction in Java and verified that it is capable of processing large numbers of bidding operations per second. Such speed is important in practice, since it is desirable for a dynamic auction to compute and publish updates to the pricing and allocation in real time.

8. REFERENCES

- [1] L. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, 2004.
- [2] L. Ausubel. An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3):602–629, 2006.
- [3] S. Bickchandani and J. Ostroy. Ascending price Vickrey auctions. *Games and Economic Behavior*, 55(2):215–241, 2006.
- [4] H. Clarke, E. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [5] S. de Vries, J. Schummer, and R. Vohra. On ascending Vickrey auctions for heterogeneous objects. *Journal of Economic Theory*, 132(1):95–118, 2007.
- [6] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.
- [7] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [8] F. Gul and E. Stacchetti. The English auction with differentiated commodities. *Journal of Economic Theory*, 92(1):66–95, 2000.
- [9] P. Jaramillo and V. Manjunath. The difference indifference makes in strategy-proof allocation of objects. *Social Science Research Network*, April 2011.
- [10] C. Krishnappa and C. G. Plaxton. A dynamic unit-demand auction with bid revision and sniping fees. Technical Report TR–11–30, Department of Computer Science, University of Texas at Austin, June 2011.
- [11] C. Krishnappa and C. G. Plaxton. A sealed-bid unit-demand auction with put options. In *Proceedings of the 22nd International Conference on Game Theory*, Stony Brook, NY, July 2011.
- [12] H. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 3:253–258, 1955.
- [13] D. Mishra and D. C. Parkes. Ascending price Vickrey auctions for general valuations. *Journal of Economic Theory*, 132(1):335–366, 2007.
- [14] D. Mishra and D. C. Parkes. Multi-item Vickrey-Dutch auctions. *Games and Economic Behavior*, 66(1):326–347, 2009.
- [15] H. Rothkopf, M. A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1999.
- [16] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.
- [17] K. Steiglitz. *Snipers, Shills & Sharks: eBay and Human Behavior*. Princeton University Press, Princeton, New Jersey, 2007.
- [18] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- [19] L. Walras. *Elements of Pure Economics*. Allen and Unwin, 1954.